Understanding Financial Market Behavior through Empirical Game-Theoretic Analysis

by

Erik Brinkman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Computer Science and Engineering)
in the University of Michigan
2018

Doctoral Committee:

Professor Michael Wellman, Chair
Assistant Professor Jacob Abernethy, Georgia Tech
Professor Scott Page
Professor Uday Rajan
Assistant Research Scientist Arunesh Sinha

"And there never was an apple, in Adam's opinion,
that wasn't worth the trouble you got into for eating it."

—Neil Gaiman & Terry Pratchett, *Good Omens*

Erik Brinkman
erik.brinkman@umich.edu
ORCID iD: 0000-0002-6712-3510

To my mother, father, and brother

ACKNOWLEDGMENTS

This thesis would not have been possible without the guidance of my advisor Michael Wellman. I had no idea what I was looking for when I started the Ph.D. program, but in Mike I found a role model, a guide, and a collegue. He fulfilled each of those roles better than I could have ever imagined. It is not an exageration to state that the majority of my development over the past six years can be attributed to Mike's influence. His constant drive for exelence has been a true source of inspiration. I am grateful for all of the time I have had under his tutalidge, and hope that this is only the beginning of our relationship.

I also owe a great debt of gratitude to my entire commitee. I will never forget the times I spent in Lab discussing future research ideas with Scott Page, the fruitful sunday morning meetings at the Ross Starbucks with Uday Rajan, the celebratory and enlightening outings with Jake Abernethy, or the random productive interractions with Arunesh Sinha in the office. They not only left a positive impression on my thesis, but have encouraged me to grow in ways I never expected. I would be remiss if I didn't mention the mentorship of Grant Schoenbeck, who behaved like an unofficial committee member.

I am grateful to all the CSE administrative staff who have helped me get here. Their contributions have helped me survive the dauting thesis process, but they have also brightened countless days with their selfless help on the smallest of chores. I specifically want to call out the kindness and assistance of Dawn Freysinger, Kimberly Mann, Rita Rendell, Yolonda Coleman, Ashley Andreae, and Kelly Gratz.

I have been fortunate to count so many bright minds as my peers and collaborators in the Strategic Reasoning Group. The continued friendship of Elaine Wah, Travis Martin, and Bartley Tablante will be one of the most positive outcomes of the Ph.D. program. I have also received countless advice and assistance from Ben Cassell, Bryce Wiedenbeck, Mason Wright, Frank Cheng, Xintong Wang, Zaina Hamid, Megan Shearer, and Thanh Nguyen.

I would not have made it this far if not for the enduring friendships that been there to support and encourage me along the way. I want to especially thank Ari

Koben, who's had my back since we first applied to grad school, and Rob Goeddel, whom I could always come bother with the most random of questions. I also want to acknowledge the many kindred spirits of the computer science department: Nilmini Abeyratne, Shruti Padmanabha, Lauren Hinkle, James Kirk, Zach Musgrave, Sanae Rosen, Lizzie Mamantov, Torrie Edwards, Andrew Richardson, Mike Cieslak, Parinazi Ardabili, and Wendy Shang. A handful of Ann Arborites and friends made during internships have continued to advocate for my success. I want to especially thank Hannah Imlay, Joey Capparella, Brittany Johnson, Gabby Melamed, Gina Sultan, Tammi Cooks, Kevin Owen, Liam Casey, Lauren Dello Russo, Shoshana Shapiro, Summer Xu, Shilpa Kobren, Agatha Radberger, and Moshe Barach.

Finally, I want to thank my parents for their unwavering support and devotion. Anything I have and will ever achieve is because they are willing to believe so strongly in what I am capable of and will sacrifice anything to make sure I'm prepared for what lies ahead.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGMENTS iii

LIST OF FIGURES vii

LIST OF TABLES ix

ABSTRACT x

CHAPTER

I. Introduction 1

II. Financial Market Simulator 5

2.1 Valuation Model 6
2.2 Background Trading Strategies 8

III. Empirical Game-Theoretic Analysis 10

3.1 Strategy Selection 12
3.2 Profile Search 12
3.3 Player Reduction 16
3.4 Bootstrapping Regret 16

IV. Shading and Efficiency in Limit-Order Markets 17

4.1 Introduction 17
4.2 Stylized Examples 18
4.3 Prior Work 23
4.4 Replication of Zhan and Friedman 25
4.5 Financial Market Environment 27
4.6 Financial Market Equilibrium Analysis 32

4.7 Conclusions 33

V. Empirical Mechanism Design for Optimizing Clearing Interval in Frequent Call Markets 38

5.1 Introduction 38
5.2 Related Work 40
5.3 Market Model 42
5.4 Analysis 44
5.5 Conclusion 53

VI. Trend-Following Trading Strategies and Financial Market Stability 55

6.1 Introduction 55
6.2 Related Work 58
6.3 Financial Market Simulator 59
6.4 Market Shocks on Equilibrium Behavior 65
6.5 Effects in Alternate Environments 70
6.6 Anticipating Market Shocks 72
6.7 Conclusion 79

VII. Conclusion 80

BIBLIOGRAPHY 83

LIST OF FIGURES

Figure

3.1	Flow chart illustrating the profile search algorithm	15
4.1	How shading can improve efficiency in the scenario of Example 1	19
4.2	Moderate shading reduces inefficiency from suboptimal trades (EM-inefficiency) without introducing significant inefficiency from missed optimal trades (IM inefficiency)	25
4.3	Efficiencies of all found equilibria compared to truth-telling in every scenario	27
4.4	Uniform shading analysis of the financial market model	32
4.5	Fraction of optimal surplus for truth-telling agents (yellow), and the efficiencies in equilibrium (blue) for each of the environments	34
5.1	Call market efficiency in the $-N\uparrow T\downarrow A$ environment as a function of clearing interval when agents are truthful	46
5.2	Equilibrium efficiency as a function of clearing interval in six environments with low adverse selection	47
5.3	Equilibrium price discovery as a function of clearing interval in low adverse-selection environments	49
5.4	Equilibrium median spreads as a function of clearing interval in low adverse-selection environments	50
5.5	How equilibrium statistics change when agents have access to call market specific strategies in environment $-N\uparrow T\downarrow A$	51
5.6	Equilibrium efficiencies in two environments with high adverse selection	52
5.7	Equilibrium price discovery with high adverse selection	52
5.8	Equilibrium spread with high adverse selection	53
6.1	Effect of market shocks in the absence and presence of price-conditional strategies	62
6.2	Price trends correlate with mispricing of the fundamental value	63
6.3	Trend following exacerbates the market's response to a price shock	64
6.4	Trend followers slightly reduce volatility in equilibrium	68
6.5	Trend followers slightly increase mispricing in equilibrium	68
6.6	Trend followers amplify the decline in price in the presence of large market shocks	69
6.7	Trend followers improve the mispricing (RMSD) in response to market shocks	70

6.8 Trend followers in markets where agents do not use transaction information produce instantaneous crashes, where price drops significantly, but recovers just as quickly 70
6.9 Reducing fundamental shock variance, one part of the friction that trend followers exploit, increases market stability 71
6.10 Reducing observation variance removes incentives to trend follow, increasing market stability 72
6.11 Mixture equilibria where multiple equilibria exist 76
6.12 Mixture equilibria where equilibria are discontinuous 76
6.13 Maximum fundamental difference with and without market shocks as the equilibrium market shock probability changes 78
6.14 Mispricing (RMSD) as the probability of anticipated shock changes 78

LIST OF TABLES

Table

4.1 Equilibria found in CDA games, by Zhan and Friedman (2007) or my replication . . . 28
4.2 Strategies considered for equilibrium analysis . . . 30
4.3 Baseline environment parameters, by setting . . . 31
4.4 Equilibria found in multi-position financial market environments . . . 35
4.5 Equilibria found in single-position financial market environments . . . 36
5.1 Strategies considered for equilibrium analysis . . . 44
5.2 Additional strategies considered for equilibrium analysis with call market strategies . . . 44
6.1 Default Environment Parameters . . . 59
6.2 Background agent strategies considered in systematic equilibrium search . . . 65
6.3 Larger set of background agent strategies considered for potential deviation from found equilibria . . . 66
6.4 Trend follower strategies considered in systematic equilibrium search 66
6.5 Larger set of trend follower strategies considered for potential deviation from found equilibria . . . 66
6.6 Empirical and bootstrap regrets for equilibria with and without trend followers in five environment settings . . . 67

ABSTRACT

Financial market activity is increasingly controlled by algorithms, interacting through electronic markets. Unprecedented information response times, autonomous operation, use of machine learning and other adaptive techniques, and ability to proliferate novel strategies at scale are all reasons to question whether algorithmic trading may produce dynamic behavior qualitatively different from what arises in trading under direct human control. Given the high level of competition between trading firms and the significant financial incentives to trading, it is desirable to understand the effect incentives have on the behavior of agents in financial markets. One natural way to analyze this effect is through the economic concept of a Nash equilibrium, a behavior profile of every agent such that no individual stands to gain by doing something different.

Some of the incentives traders face arise from the complexities of modern market structure. Recent studies have turned to agent-based modeling as a way to capture behavioral response to this structure. Agent-based modeling is a simulation paradigm that allows studying the interaction of agents in a simulated environment, and it has been used to model various aspects of financial market structure. This thesis builds on recent agent-based models of financial markets by imposing agent rationality and studying the models *in equilibrium*.

I use empirical game-theoretic analysis, a methodology for computing approximately rational behavior in agent-based models, to investigate three important aspects of market structure. First, I evaluate the impact of strategic bid shading on agent welfare. Bid shading is when agents demand better prices, lower if they are buying or higher if they are selling, and is typically associated with lower social welfare. My results indicate that in many market environments, strategic bid shading actually improves social welfare, even with some of the complexities of financial markets. Next, I investigate the optimal clearing interval for a proposed market mechanism, the frequent call market. There is significant evidence to support the idea that traders will benefit from trading in a frequent call market over standard continuous double auction markets. My results confirm this statement for a wide variety of market

settings, but I also find a few circumstances, particularly when large informational advantages exist, or when markets are thin, that call markets consistently hurt welfare, independent of frequency. I conclude with an investigation on the effect of trend following on market stability. Here I find that the presence of trend followers alters a market's response to shock. In the absence of trend followers, shocks are small but have a long recovery. When trend followers are present, they alter background trader behavior resulting in more severe shocks that recover much more quickly. I also develop a novel method to efficiently evaluate the effect of shock anticipation on equilibrium. While anticipation of shocks does make markets more stable, trend followers continue to be profitable.

CHAPTER I

Introduction

Over the past couple of decades the landscape of U.S. financial trading has changed significantly. Automated trading algorithms accounted for over half of all U.S. equities trading in 2017, up from a little over a quarter in 2006 (Meyer, Bullock, and Rennison 2018). Much of this activity—known as *high-frequency trading* (HFT)—is from algorithms that have incredibly fast access to market information and react near the fundamental limits of communication, These developments have evolved our financial system from one where changes took place on the order of minutes and every event was the result of human decision making, to one where every market in the world is connected by low-latency, autonomous, independent algorithms. These algorithms respond to both internal and external information, creating feedback and market dynamics that are hazily understood at best.

This shift in market dynamics and its potential effect on participants hasn't gone unnoticed. Many academics and financial insiders have argued about whether HFT is beneficial or detrimental to modern markets. Evidence suggests that HFT improves liquidity and enhances the informativeness of quotes (Hendershott, Jones, and Menkveld 2011), but it also fuels wasted investment in low-latency communication and reduces the welfare of other market participants (Budish, Cramton, and Shim 2015; Wah and Wellman 2013). While the two negative studies cited present different causes for welfare loss, each cause can be traced back to one underlying feature of modern markets: trades are matched as soon as they arrive. This instantaneous matching results in a winner-takes-all system for the fastest participant. A proposed solution to restore welfare is to switch from continuous time markets to discrete time call markets with frequent periodic clearing, which reduce the benefits from incremental speed advantages. While there has been some research surrounding frequent call markets, including their effects on HFT and how they could comply with market regulation (Budish, Cramton, and Shim 2014), most of the existing studies that suggest a choice of frequency do so with little evidence.

Perhaps the most obvious and pressing effect of these financial market changes is the rise in catastrophic financial events. Some such events are simply the result of negligence, like the 2012 Knight Capital incident, where a software bug resulted in the loss of over $440 Million, and a subsequent indictment by the SEC (SEC 2013). But some, like the 2010 Flash Crash, where major equity indices fell 5% in minutes before rebounding almost as quickly, were still being actively investigated as recently as last year (Kirilenko et al. 2017). The Flash Crash has been attributed to multiple causes including a large misplaced order and alleged market manipulation (Popper and Anderson 2015), although the latter is widely disputed as a cause of the Flash Crash. The lengthy analysis and general uncertainty about factors contributing to the Flash Crash underscores the difficulty in understanding the cause of catastrophic market events.

In order to explore the issues brought on by automated trading, I use a methodology that I believe exposes aspects of modern financial markets that are difficult to express via alternative methods. There are three common ways with which scientists study financial markets, *1)* using analytic models and classical game theory to predict the behavior of agents, *2)* using historic data and financial analysis to extrapolate future market behavior, *3)* and using agent-based models to ascribe behavior to market participants and study market outcomes predicated on that behavior.

Analytic models have proven to be quite useful; they sometimes allow characterization of the space of equilibria beyond simple enumeration, and they offer a precise characterization of effects within their models (Budish, Cramton, and Shim 2015; Du and Zhu 2017; Satterthwaite and Williams 1989). However, analytic models are often simplified to make them more amenable to analysis techniques. For example, I am unaware of any analytic model that incorporates agent reentry in a continuous double auction (CDA), an important detail of the models presented in this thesis. This is not to say that analytic models are not insightful, but they often remove complexity by design, which is a problem if the complexity is a relevant feature.

A second approach, using empirical analysis of historic financial data, can remedy this problem as it reflects whatever complex interactions generate it. However, historic data is observational, and it can only be used to answer observational questions, like “are market crashes frequently preceded by a large sell order?” What they cannot directly address are counterfactual questions, like “will a novel market intervention prevent large sell orders from triggering market crashes?”

The third technique lies somewhere between these first two. Agent-based models simulate the interactions between autonomous agents. They can account for the complexity lacking in analytic models, by expressing any desired complexity in simulation. Agent-based models also provide a means to answer counterfactual questions. Once created, the model can run under many environments to investigate how a new environment changes qualitative or quantitative behavior. These models still have one large flaw: the behavior of agents is entirely prescribed by the model designer. Agents change their behavior only in explicitly designed ways. Often, these models are calibrated to market data, but the rules agents follow remain fixed.

The technique employed in this thesis, known as Empirical Game-Theoretic Analysis (EGTA) (Wellman 2006), improves on standard agent-based models by allowing agents to *strategically* react. Instead of prescribing agent behavior in an environment, EGTA uses parameterized strategies and agent utilities to induce a normal-form game from the agent-based model. From this game, any game-theoretic solution concept can be used to determine agent behavior. In the market settings investigated in this thesis we solve for Nash equilibria, where agents are mutually best responding. This makes aggregate agent behavior less dependent on the design of individual agent strategies because agents will adopt different strategies depending on the utility they achieve for adopting them. The result is agent behavior that is more dependent on the utilities and solution concept, which in turn reduces the potential for experimenter bias in strategy design. In addition, Nash equilibration with even a set of simple parameterized strategies often makes agents competitive with more complicated strategies. The general form of the agent-based model I employ is detailed in Chapter II, and the details of the EGTA analysis are covered in Chapter III.

This dissertation analyzes three different case studies of agent behavior in financial markets. I investigate when strategic bid shading improves agent welfare over truth telling in Chapter IV. Bid shading is when agents submit orders that guarantee some level of profit conditioned on transaction. This reduces the chance of an order transacting, which can lower overall welfare, but is often beneficial to shading agent. My results indicate that even with the complexities of financial markets, strategic bid shading improves welfare in most environments, mirroring prior work with simple CDAs (Zhan and Friedman 2007). In Chapter V I use empirical mechanism design to evaluate how the length of the clearing interval in frequent call markets affects agent welfare. I find that in many environments a sufficiently short clearing interval provides a benefit to agent welfare. However, in markets that are thin and have a large amount of adverse selection, frequent call markets hurt welfare. I conclude with a study on how trend followers, agents that extrapolate value from price trends, affect market stability in Chapter VI. My results indicate that trend followers alter the dynamics of how a market responds to external shocks. Without trend followers, background agents have incentive to use price information in their value estimate, causing recovery after a shock to be slow as agents' beliefs reset. Trend followers crowd out background agents, removing their incentive for using price information. As a result, market shocks with trend followers are much more severe, amplified by the trend followers, but recover much more quickly because agent beliefs do not need to reset. The models and methodologies underpinning these results are robust and extensible to further forms of analysis predicated on agent rationality. They build on previous work in this area and continue to demonstrate that empirical microeconomic models of financial markets produce useful insight into market microstructure.

CHAPTER II

Financial Market Simulator

My experiments employ a configurable financial market simulator, originally developed by Elaine Wah and since edited by student researchers in the Strategic Reasoning Group, including many contributions by myself. The simulator in some form has been used in many experiments covering topics as diverse as strategic market choice (Wah, Wright, and Wellman 2017) and spoofing (Wang and Wellman 2017). Actors in financial markets hold multiple units of a security which they both buy and sell. They frequently do not have perfect information about the security's value, and they often trade with more informed agents. This simulator was designed to capture all of these qualitative aspects in order to accurately model the agent interactions and incentives in modern financial markets.

Almost all modern financial markets use the same basic mechanism. Agents interact with the market by placing limit orders representing the maximum (minimum) price an agent is willing to buy (sell) a single unit of the security. If at any point two agents are willing to transact—one agent's maximum price to buy a unit is greater than or equal to another agent's minimum price to sell a unit—the orders match, and the agents trade at the price of the incumbent order. This mechanism is also known as a continuous double auction (CDA).

For simplicity, I model a single security traded in a CDA market. Prices and time are fine-grained but discrete, and agent interactions are strictly ordered. If two agents arrive at the market at the same time t they act in a random order; one will always interact with the market before the other. Therefore, even though time is discrete, agent interaction with the market is continuous. Simulations run for a finite horizon T. The market maintains and reveals price quotes reflecting the best outstanding orders at the current time as well as past transaction prices; other bids in the market are not visible to agents.

2.1 Valuation Model

Each agent has an individual valuation for the security comprised of private and common components. The private component provides agents a reason to trade. There are many reasons for why an agent would have a private valuation for security, such as liquidity demands or hedging with a correlated asset. This simulation simply models a general private value as a way to abstract an agent's desire to trade. The common component serves to model the resale value of a security and to provide a source of adverse selection, described shortly. The two components are added together to compute an agent's utility.

The common component is modeled as a stochastic *fundamental value*. Let f_t denote the fundamental value for the security at time t. The fundamental is generated by a mean-reverting stochastic process, the discrete analogue of the Ornstein-Uhlenbeck process (Doob 1942):

$$f_t = r\bar{f} + (1-r)f_{t-1} + s_t, \qquad f_0 = \bar{f}, \qquad s_t \sim \mathcal{N}\left(0, \sigma_s^2\right). \tag{2.1}$$

s_t is a random Gaussian shock at time t, with variance σ_s^2—hereafter referred to as the fundamental shock variance. The fundamental shock variance controls the volatility of the fundamental value. Parameter $r \in [0,1]$ specifies the tendency by which the fundamental reverts back to the mean $\bar{f}$; $r = 0$ corresponds to a martingale Gaussian fundamental.

The time-varying fundamental presents agents with an issue of *adverse selection* (Akerlof 1970)—trading with agents who have better information—as standing orders reflect outdated information from the time submitted. If the fundamental shifts significantly, subsequently arriving agents are more likely to transact with orders on the side opposite the direction of change. That is, a positive price shock will tend to trigger transactions with stale sell orders, and negative price shocks with stale buys. Say there are two agents B and S, where B has a \$10 private value to buy and S has a \$10 private value to sell. Say Agent B observes the fundamental at \$100 and submits a buy order at \$100 attempting to make \$10 from their private value. After submitting the order, a negative price shock occurs, and the fundamental drops to \$50. Then Agent S arrives and observes the fundamental. Agent S is going to take the resting order for \$60 of profit, leaving Agent B with a \$40 loss. A similarly sized positive shock does not compensate for Agent B's loss, as Agent S will choose not to transact instead. Strategic agents who leave orders in the market must ask for additional profit, beyond their unbiased estimate of the marginal value of transaction, to account for these events. This strategic bid adjustment is known as *shading*.

The degree of adverse selection in this model depends on both the fundamental shock variance σ_s^2 and the degree of mean reversion r. As elaborated in the previous paragraph, an agent risks trading with inferior information if they leave a resting order in the market and there is a subsequent shift in the fundamental value. If the fundamental shock variance increases, the probability of large fundamental shifts also increases, exposing agents to more adverse selection. Higher values of mean reversion damp early fundamental variations, making their effect on the final fundamental value small, thus a higher mean reversion exposes agents to less adverse selection. To illustrate this point, imagine the mean reversion, r, equals 1. After every shock, the fundamental returns immediately to the mean; only agents arriving at the final time step have an information advantage.

The private component of agent i's valuation is a vector $\boldsymbol{\theta}_i$ containing the agent-specific marginal utility for acquiring one more unit, relative to the fundamental value—similar to the model of Goettler, Parlour, and Rajan (2009). The vector is of size $2q_{\max}$, where $q_{\max} > 0$ is the maximum number of units the agent can be long or short at any time. Element $\theta_i^q, q \in [-q_{\max}, q_{\max}) \cap \mathbb{Z}$ is the incremental benefit, over the fundamental, to agent i for gaining one unit of the security given current position q, where positive q indicates a long position.

$\boldsymbol{\theta}_i$ is generated from a set of $2q_{\max}$ values drawn independently from a Gaussian distribution. Let $\hat{\theta} \sim \mathcal{N}\left(0, \sigma_{pv}^2\right)$ denote one of these drawn values. To ensure that the valuation reflects diminishing marginal utility, the $\hat{\theta}$ values are sorted in descending order and assigned to $\boldsymbol{\theta}_i^q$ respectively.

Agent i's incremental surplus for a trade is based on its position q before the trade, the value of the fundamental at the end of the trading horizon T, and the transaction price p:

$$\text{surplus} = \begin{cases} f_T - p + \theta_i^q & \text{if buying 1 unit} \\ p - f_T - \theta_i^{q-1} & \text{if selling 1 unit.} \end{cases}$$

An agent's total surplus is the sum of the agent's surplus over all transactions. Since the price and fundamental terms cancel out in exchange, the total surplus achieved when agent B buys from agent S is $\theta_B^{q_B} - \theta_S^{q_S - 1}$, where q_i denotes the pre-trade position of agent i.

2.2 Background Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs (R. Das et al. 2001; Friedman and Rust 1993; Wellman 2011). In this thesis, I primarily consider trading strategies that are variants of the *Zero Intelligence* (ZI) family (Gode and Sunder 1993). The ZI strategy is exceedingly simple, but often employed in agent-based study of financial markets (Cason and Friedman 1996; LeBaron 2006), including recent AI studies of market making (Chakraborty, S. Das, and Peabody 2015; Wah, Wright, and Wellman 2017), because they have been found to generate realistic patterns of market behavior (Farmer, Patelli, and Zovko 2005). Though ZI agent instances are typically outperformed by more sophisticated alternatives (Tesauro and R. Das 2001; Vytelingum, Cliff, and Jennings 2008), using game-theoretic selection to set ZI strategy parameters can produce highly competitive behavior for a given market environment (Wright and Wellman 2018). Game-theoretic equilibration competitively tunes naive ZI agent parameters, a result achieved through adaptation and evolutionary search by the *ZI Plus* agents of Cliff (2009).

In this market model, agents get information about the fundamental according to an independent geometric process with probability λ of arriving and observing the fundamental (f_t) at any specific time. Upon observing the fundamental, agents also observe the price quote, withdraw any outstanding orders, and have the opportunity to submit new ones. If multiple agents get information at the same time step, they act in a random order. Agents are assigned on each arrival to either buy or sell, with equal probability, and accordingly submit an order to buy or sell a single unit. This style of order submission is representative of studies employing the standard ZI strategy. Agents may trade any number of times, as long as their net positions do not exceed their maximum position ($q_{\max}$) (either long or short).

A ZI agent assesses its expected valuation at the time of market entry t, using an estimate $\hat{f}_t^T$ of the terminal fundamental f_T. The estimate is based on the current fundamental, f_t, adjusted to account for mean reversion:

$$\hat{f}_t^T = (1-\rho)\,\bar{f} + \rho f_t, \qquad\qquad \rho = (1-r)^{T-t}. \tag{2.2}$$

The ZI agent then submits a bid *shaded* from this estimate by a random offset—the amount of expected surplus it demands from the trade. The amount of shading is drawn uniformly from the range $\left[\underline{d}, \overline{d}\right]$. $\underline{d}$ and $\overline{d}$ are strategic parameters that agents can choose.

Background agents use an extended form of the ZI strategy (Wah, Wright, and

Wellman 2017) that includes a strategic threshold parameter $\eta \in [0,1]$, whereby if the agent could achieve a fraction η of its requested surplus at the current price quote, it simply takes that quote rather than posting a limit order to the book. Setting $\eta = 1$ is equivalent to the strategy without employing the threshold. Settings of $\eta < 1$ are often highly advantageous in my simulation environments, suggesting that providing even this simple ability to condition on price quote is an important feature in CDA trading strategy.

The threshold parameter can also be used as a means to submit so-called "fill-or-kill" orders, designed to trade immediately or not at all. Consider the strategy where $\underline{d} = \overline{d} = DZ$ and $\eta = Z^{-1}$ for very large Z. This strategy takes the outstanding order at the quote—if a surplus of at least D is available—or else posts an order at such an unattractive price that it will never transact. Interestingly, a similar strategy appears in empirical equilibrium in several of the analyzed environments.[1]

This model, as presented, is still missing key features necessary to analyze the scenarios in later chapters. In Chapter IV, I extend the background order submission strategy by allowing agents to place simultaneous buy and sell orders. This modified strategy combined with with a maximum position of one makes it so agents can achieve their competitive surplus position in one arrival—removing trader urgency as motivation to shade less. Chapter V extends the model by adding call markets, markets that clear at periodic intervals instead of at order submission, and agent strategies tailored to call markets. In Chapter VI I add imperfect fundamental observations and trend-following agents to the model. While each chapter represents a unique contribution to market modeling, the incremental scope of each extension speaks to the robustness of this base model.

[1]Since the fill-or-kill strategy never leaves reasonable orders, in self-play it never trades. It can however be profitable in mixtures with other strategies.

CHAPTER III

Empirical Game-Theoretic Analysis

Traditional game theory focuses on games with full knowledge of of the payoffs for every strategy profile, the assignment all players to strategies. This information typically comes from enumerated payoffs in small games, or larger games with an entirely analytic specification. Many sufficiently complex game models may not yield a tractable analytic form or may otherwise be too large for distinct enumeration of payoff values, yet may be relatively straightforward to simulate and obtain samples from the payoff distribution of a profile. These games are generally referred to as simulation-based games (Vorobeychik and Wellman 2008).

The standard procedure for analyzing simulation-based games can be broken down into four parts:

1. Exploit any symmetry in the game by assigning the players to *roles*. Roles are groupings of agents that are symmetric. For two agents to share a role, each agent must have access to the same strategies, be ex-ante identical, lack coordination, and be anonymous (switching the agents has no effect on the rest of the game). I will refer to a game with roles as *role symmetric*. A fully symmetric game is role symmetric with a single role; any arbitrary game is role symmetric with a unique role for every agent. A role-symmetric mixed-strategy equilibrium is an equilibrium of a role-symmetric game, where every player within a role independently chooses a strategy from the same distribution. Every finite role-symmetric game is guaranteed to have at least one role-symmetric mixed-strategy Nash equilibrium.

2. If the game has a *large* strategy set select some finite subset of the strategies to covert the game to normal form. What defines a large strategy set varies by computation time of the simulator and available resources, but two common

game types that require subset selection are extensive form games, where the number of strategies scales exponential with the depth, and continuous games, which must have a finite subset of strategies selected.

3. Define how to evaluate profiles. A profile is evaluated if the payoffs to each player are known with reasonable certainty. If simulator payoffs are deterministic then only one run of the simulator is necessary to evaluate the payoffs, but if simulator payoffs are samples from a distribution, then some sufficient number of runs is necessary to get a low variance estimate of the expected payoffs for that profile. Selecting a number of profiles in advance is not strictly necessary, nor is it necessarily optimal, but adaptively evaluating profiles is beyond the scope of this thesis.

4. Evaluate profiles in such a way to construct role-symmetric normal-form empirical games, and use numeric techniques to identify solutions in these created games. Simulation-based games often have too many distinct profiles for complete profile evaluation. I use the techniques in Sections 3.2 and 3.3 to evaluate fewer profiles. In this thesis, I aim to identify approximate role-symmetric mixed-strategy Nash equilibria. An approximate equilibrium has regret bounded below some small threshold, where the regret is the maximum an agent can gain by unilaterally deviating to any strategy. Approximate solutions are, to some extent, necessary as I use numerical methods from game data to compute them and the results will be subject to some numerical instability. The regret also serves as a measure of plausibility; if the gain from deviating to a new strategy is small, agents are less likely to perceive a benefit from deviating.

This entire process, as well as related techniques for analyzing the empirical data from simulation-base games are known under the umbrella of *empirical game-theoretic analysis* (EGTA) (Wellman 2006).

The model presented in Chapter II and its derivatives presented later in this thesis have no known analytic solution, but are readily amenable to simulation, making them ideal for use with EGTA. In each of the remaining chapters, I use the techniques presented here to identify role-symmetric mixed-strategy equilibria in the relevant game models.

3.1 Strategy Selection

To make the games amenable to normal-form game analysis, I choose a finite set of individual strategies from the entire strategy space to define the game of interest. Such a choice inevitably loses guarantees with respect the true strategy space, but my goal is to minimize the chance that equilibria in the game with a finite set of strategies have low regret in the true game. To this end, each analyzed strategy set was chosen via an iterative process inspired by Schvartzman and Wellman (2009). Starting from an original set of candidate strategies, I identify role-symmetric equilibiria within that set. I then attempt to find beneficial deviations outside the strategy set, and upon finding one add it to the finite set considered for equilibria. This search gives better confidence that the reported equilibria have low regret in the full game.

3.2 Profile Search

Even with a finite number of strategies, the number of distinct profiles that would need to be evaluated to induce payoffs for a complete normal-form game is astronomical, significantly larger than the number of stars in the universe.[1] I employ a heuristic profile search that biases towards low-support equilibria—equilibria where a small set of strategies are played with nonzero probability—to reduce the number of profiles that need to be evaluated. By evaluating only a small number of strategies played simultaneously I hope to confirm equilibria while only having to evaluate a small number of profiles. I adopt a set of heuristic criteria that promote thorough coverage of the profile space, in order to avoid being too unduly biased by equilibria found early in the search process. This search process was heavily inspired by the work of Jordan, Schvartzman, and Wellman (2010).

The search criteria rely on *exploring* games where agents are restricted to playing only a small subset of the strategies, referred to as *restricted games*. I consider a restricted game *explored* if I have *(i)* evaluated every profile in the restricted game, *(ii)* found at least one candidate equilibrium in the restricted game, and *(iii)* for each candidate equilibrium, evaluated all of the one-player deviations to strategies outside the restricted game . A *candidate equilibrium* is an equilibrium of a restricted game, but whose deviations to strategies in the unrestricted game may be unevaluated, thus it may or may not be an equilibrium in the full game. The stopping criteria are as follows:

[1]There are more 10^{57} profiles in the largest games I study and an estimated 10^{24} stars in the universe (*How many stars are there in the universe?* 2004).

Criterion 1. *I have found at least one role-symmetric Nash equilibrium.*

Criterion 2. *I have explored every restricted game with a single strategy per role (i.e., every strategy in self-play).*

Criterion 3. *For every candidate equilibrium found in an explored restricted game, I have explored the restricted game formed by adding the best-response strategy to the support of that candidate, providing that the number of strategies in the restricted game is below a threshold.*

To meet these criteria, I iteratively apply EGTA using a process adapted from the inner loop of the procedure defined by Wellman, Kim, and Duong (2013). Pseudocode for the profile search process is listed as Algorithm 1 and Figure 3.1 features a diagram roughly describing the process.[2]

The search process starts with a set of restricted games $\mathcal{R}$ that are required to be explored per the criteria. We adopt the notation that a restricted game is the set of allowed strategies. This set is initialized to the restricted games comprising exactly one strategy per role.

For each restricted game G in this set, evaluate all of the profiles in the restricted game and then compute Nash equilibria on the empirical game. Finding a Nash equilibrium is PPAD Complete (Daskalakis, Goldberg, and Papadimitriou 2009), but often a couple of polynomial-time incomplete algorithms—such as regret minimization and replicator dynamics—are sufficient to find at least one equilibrium. These algorithms are incomplete because they may fail to find an equilibrium in the restricted game. In the event that that happens, I run simplicial subdivision for fixed points to guarantee an equilibrium is found.

For every candidate equilibrium Q found in a restricted game, I evaluate every deviating profile to the unrestricted game. If none of the deviations is beneficial, then the candidate is a confirmed equilibrium and saved in set $\mathcal{C}$. Otherwise, if the restricted game has few enough strategies, add the best response to the restricted strategy set and explore the new restricted game. Other beneficial-deviation restricted games are kept as backups in set $\mathcal{B}$ to explore in case every required restricted game has been explored but no equilibria were found. In the unlikely event that no backup restricted games exist, I take a random unexplored restricted game with the lowest support and explore it.

[2] My implementation of this procedure is available at https://github.com/egtaonline/quiesce. It is integrated with EGTA Online and fully automatic.

Algorithm 1 Profile Search

define br ▷ Returns the best-response strategy to a restricted game
define brs ▷ Returns all beneficial response strategies to a restricted game
define ξ ▷ Restricted game size threshold
define $\mathscr{P}$ ▷ The power set
define $\mathcal{S}$ ▷ The set of strategies

```
C ← ∅                                         ▷ Confirmed equilibria
R ← {{s} | s ∈ S}                             ▷ Required restricted games
E ← ∅                                         ▷ Explored restricted games
B ← ∅                                         ▷ Backup restricted games
while C = ∅ or R \ E ≠ ∅ do
    if R \ E ≠ ∅ then
        G ∈ R \ E                             ▷ Required game
    else if B \ E ≠ ∅ then
        G ∈ B \ E                             ▷ Backup game
    else
        G ∈ argmin_{G' ∈ P(S) \ E} |G'|       ▷ Minimum size unexplored game
    end if
    E ← E ∪ {G}
    Evaluate all profiles in G
    for all equilibria Q of G found do
        Evaluate all deviating profiles from Q
        if brs(Q) = ∅ then
            C ← C ∪ {Q}                       ▷ Add confirmed equilibrium
        else if |G| ≤ ξ then
            R ← R ∪ {G ∪ br(Q)}               ▷ Explore best response
            B ← B ∪ {G ∪ {s} | s ∈ brs(Q) \ {br(Q)}}
        else
            B ← B ∪ {G ∪ {s} | s ∈ brs(Q)}
        end if
    end for
end while
```

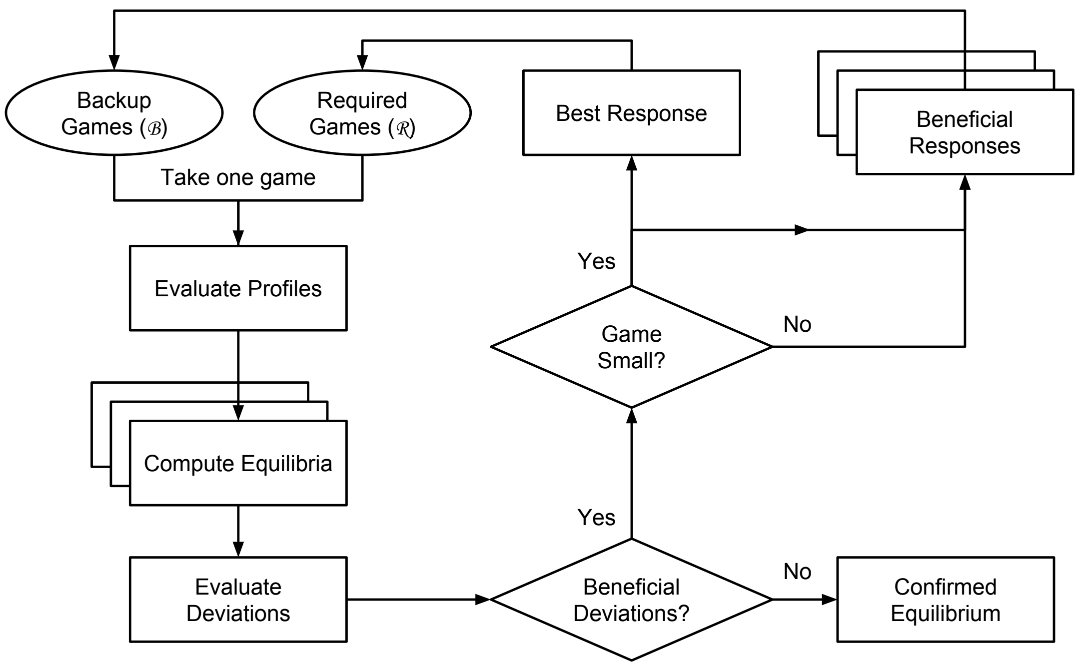

Figure 3.1: Flow chart illustrating the profile search algorithm. Required games starts with all of the restricted games with only one allowed strategy per role. For each restricted game, profiles are evaluated to compute equilibria along with their deviations. If there are no beneficial deviations, we have confirmed an equilibrium. Otherwise, if the game is small enough we require the best response be evaluated. Finally add any beneficial deviations to backup games in case we've explored all required games but found no equilibria.

3.3 Player Reduction

The number of profiles necessary to evaluate in low-support restricted games is still infeasible if there are a large number of players. I employ an aggregation technique known as *deviation-preserving reduction* (DPR) (Wiedenbeck and Wellman 2012) to approximate games with many players as reduced games with a smaller number of players. DPR is a heuristic that constructs a reduced normal-form game with payoffs from the unreduced game. It selects payoffs such that the regret of symmetric pure profiles is preserved and the regret of symmetric mixed profiles is approximated. The approximation for symmetric mixed profiles typically works well if the payoff impact of single other agents is small. The number of profiles DPR evaluates is proportional to the number of profiles in the reduced game. While other methods to reduce profile evaluation in games with many players exist (Wiedenbeck, Yang, and Wellman 2018), their usage in evaluating large games is still experimental.

3.4 Bootstrapping Regret

The combination of all of these techniques allows low-support-biased discovery of role-symmetric equilibria in the DPR reduced game, but not in the full-player game. Despite knowing that our equilibria have low regret in the reduced game, it is important to estimate how plausible they are by putting confidence bounds on the regret in the full-player game.

For every confirmed equilibrium in the reduced game, I gather additional deviating payoff samples from the mixture distribution in the full-player game and use bootstrapping to compute an upper confidence bound on regret. This is a slight variation on the bootstrap regret technique of Wiedenbeck, Cassell, and Wellman (2014), and has been shown to produce calibrated confidence bounds.

The heuristics these techniques employ mean that I have no guarantees about the coverage of the whole solution space. More pragmatically, evaluating every profile from even a relatively small game might be prohibitively expensive, as the number of distinct profiles in a symmetric game grows exponentially. In spite of these difficulties, the combination of these analysis techniques allows the use of game theory to analyze strategic interaction between agents that is much richer than tractable analytic models can provide.

CHAPTER IV

Shading and Efficiency in Limit-Order Markets

4.1 Introduction

The *continuous double auction* (CDA) (Friedman 1993) is a simple and well-studied auction mechanism, ubiquitous as the mechanism implementing limit-order markets for financial trading. All major stock exchanges, as well as commodity and futures trading institutions, employ some form of CDA for the bulk of their trading activity, and have for a century or more. Given the centrality of this market institution, it is striking that some basic strategic properties remain hazily understood. Financial trading features dynamic interaction with incomplete information (both private and common value elements), in a setting where one's own bid can determine the price of the current transaction, and influence prices of subsequent trades. To date, auction theory has not successfully tackled this combination of issues in a setting approaching the richness of financial markets.

It has long been understood that CDAs can exhibit *allocative inefficiency* (Gode and Sunder 1997): at the end of trading the goods may not be held by those who value them the most. Biais, Glosten, and Spatt (2005) survey an extensive literature on reasons for inefficiency, and ways in which concentrating trading in time (as in a call market) can improve efficiency. Prior studies in my research group have also found efficiency gaps in continuous trading (Wah and Wellman 2013). Stated simply, when agent arrivals to the market are spread over time and randomly ordered, the myopic matching procedure of the CDA may produce suboptimal allocations.

In addition, strategic bidders *shade* their bids from their true values to account for the potential effect of their bids on the transaction price. In a *one-shot* double auction setting, such shading can only degrade allocation quality, from the perfect efficiency that would be achieved with truthful bidders. Strategic financial agents should also be expected to shade their bids, so it bears considering how that affects allocative efficiency in a dynamic limit-order market.

This question was previously addressed within a standard CDA model by Zhan and Friedman (2007), who found that profiles of shading (markup) strategies in a restricted form of Nash equilibrium are highly efficient, and often yield better allocations than truth-telling. I replicate key parts of this prior work and extend it in two major directions. First, I conduct a more comprehensive game-theoretic analysis, evaluating a much larger space of strategy profiles and considering mixed as well as pure strategies in my search for Nash equilibria. Second, in addition to their standard model, I also investigate a richer family of market environments (derived from the model in Chapter II) designed to capture key features of financial markets, including private and common valuation elements, significant dynamic structure, and a broader space of agent strategies. Importantly, the common value element of this market model introduces adverse selection, which adds an extra incentive for agents to shade their bids that is not present in independent private value model of previous work.

Like Zhan and Friedman, I evaluate profiles of shading strategies through simulation. Analytic solutions of CDA games are intractable beyond simple instances, and all the more out of reach for my richer class of financial trading scenarios. I conduct an extensive empirical game-theoretic analysis, to identify the direction of the effect on efficiency in the region of equilibrium trading behavior. My results confirm that the efficiency improvements of pure-strategy equilibria found by Zhan and Friedman are also exhibited by mixed-strategy equilibria in their simple model and in more complex financial markets. However, shading equilibria are not always more efficient than their truth-telling counterparts. Only with a large number of agents, meaningfully limited agent arrivals, or large adverse selection do I find outcomes produced by strategic bidders superior to the results of truth-tellers.

4.2 Stylized Examples of Shading's Effect on CDA Equilibrium Efficiency

I develop insight on the effect of strategic shading on CDA efficiency through inspection of some simple CDA bidding scenarios. I measure efficiency by the ratio of the expected total surplus of a profile to the expectation of the maximum surplus possible—also known as the competitive equilibrium (CE) surplus. As I show, equilibrium shading can either increase or decrease overall efficiency when compared

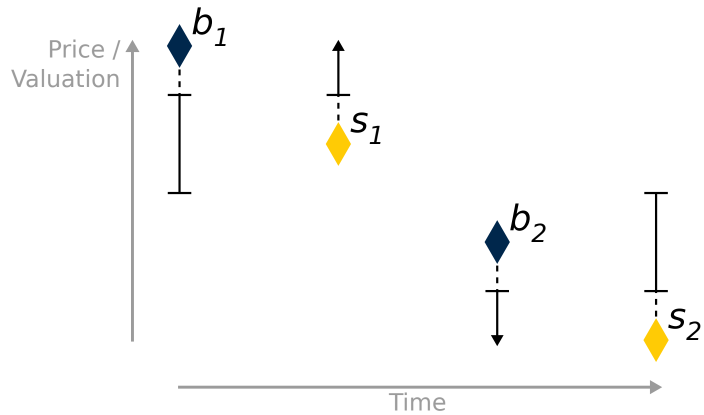

Figure 4.1: How shading can improve efficiency in the scenario of Example 1. Diamonds represent agent valuations, dashed intervals represent inefficient shading, and solid intervals represent efficient symmetric shading.

to truth-telling. In this section, I present an instance that demonstrates *how* shading can benefit efficiency, then I analyze an instance where the benefits are obtained in perfect Bayesian equilibrium (PBE), and put forth another instance where PBE shading degrades surplus.

The first example was previously employed by Wah, Wright, and Wellman (2017) to illustrate the allocative inefficiency of CDAs.

Example 1. *Consider a market with two buyers and two sellers. The buyers have private values b_1 and b_2, and sellers have private values s_1 and s_2, such that values are ordered $b_1 > s_1 > b_2 > s_2$.*

Suppose that the agents arrive at the market in order from greatest valuation to least valuation. This sequence is shown in Figure 4.1. If the agents indeed submit orders at their valuations (the diamonds), then buyer 1 trades with seller 1, yielding surplus $b_1 - s_1$, and buyer 2 trades with seller 2, yielding $b_2 - s_2$. If instead buyer 1 just trades with seller 2, the total surplus is $b_1 - s_2$, which is socially optimal. The greedy matching of the CDA in this instance executes trades that preclude efficient allocation. Whether this happens depends on how the limit orders are sequenced. With bids priced at these valuations, a random permutation of limit orders has a two-thirds probability of being suboptimal.

Suppose instead that the agents shade their bids away from their true valuations. For simplicity, they all shade symmetrically, and their valuations are uniformly spaced. If the agents shade more than $\frac{1}{2}(b_1 - s_1)$, then the inefficient trades will not happen. As long as b_1 and s_2 shade less than $\frac{1}{2}(b_1 - s_2)$, then these two still trade, and the allocation is efficient.

There is a simple explanation for why a little shading promotes efficiency in CDA markets. When there is variation in agent values, competitive equilibrium trades provide relatively high surplus to most agents involved, and are therefore tolerant to some amount of shading. In contrast, the inefficient trades generate lower surplus and can be prevented by a moderate amount of shading.

Having demonstrated that non-shading strategies can produce inefficient outcomes, and that some shading can restore efficiency, the natural next question is what happens in equilibrium? Strategic agents will clearly shade bids away from their valuations. Will the shading levels be sufficient to restore efficiency, or will they perhaps shade too much? The next example is a stylized situation where PBE is more efficient.

Example 2. *Let there be two buyers and one seller. One buyer, BL, has valuation* v_{BL} *with* $\Pr(v_{BL} = 1) = \Pr(v_{BL} = 2) = 0.5$. *The other buyer, BH, has valuation* v_{BH} *with* $\Pr(v_{BH} = 2) = \Pr(v_{BH} = 3) = 0.5$. *The sole seller, S, has valuation* $v_S = 0$. *The agents arrive at the market in a uniform random ordering, and get to observe the order book and their position in the ordering.*

In the optimal outcome BH trades with S, for expected welfare (total surplus) 2.5.

The truth-telling outcome is straightforward. If everyone bids their valuations, then when S arrives first, it will trade with the first buyer to enter the market (welfare $= \frac{1}{2}\mathbb{E}[v_{BL}] + \frac{1}{2}\mathbb{E}[v_{BH}] = 2$). When S arrives second, it will also trade with the first buyer to arrive (welfare $= 2$); when S arrives last, it will trade with the best order in the market (welfare $= 2.5$). The expected welfare under truth-telling is $\frac{1}{3}(2 + 2 + 2.5) \approx 2.167$.

To characterize strategic behavior in this example, I adopt perfect Bayesian equilibrium (Fudenberg and Tirole 1991) solution concept. A perfect Bayesian equilibrium is a solution concept of a sequential game with incomplete information where each agent updates their beliefs according to Bayes's rule and acts optimally according to those beliefs. I also assume for this construction that players break ties by accepting indifferent trades; this assumption does not affect welfare, but allows me to avoid discussing the multitude of qualitatively similar equilibria due to indifference. If S enters first, it bids 2. This offer will always trade, since at least one and possibly both of the subsequently arriving buyers have value at least this high.

If a buyer arrives first, the PBE behavior of that buyer can be described by the following cases:

1. BL arrives first, $v_{BL} = 1$. BL offers any $x < 1$.

2. BL arrives first, $v_{BL} = 2$. BL offers $\frac{4}{3}$.

3. BH arrives first, $v_{BH} = 2$. BH offers $\frac{4}{3}$.

4. BH arrives first, $v_{BH} = 3$. BH offers $\frac{4}{3}$.

If the seller arrives second, it will reject any offer less than $\frac{4}{3}$, and ask for 2, otherwise it will accept. If another buyer arrives second, it will bid ϵ over the current bid if it can profit from doing so. The last agent will take the best offer it can.

First I show that when S arrives second, it is best responding under its presumption that the third agent is a price-taker. In Case 1, the third agent is BH (and S infers this), in which case offering to sell at 2 maximizes surplus. Cases 2–4 form an information set for S, in which the third agent is BL with probability $\frac{2}{3}$. If S were to reject the first agent's bid, its optimal offer would be 2, yielding an expected surplus of $(\frac{2}{3})(\frac{1}{2})2 + (\frac{1}{3})2 = \frac{4}{3}$. So it may as well accept the first agent's bid. Given S's strategy, no first-round buyer bid less than $\frac{4}{3}$ would suffice. The only advantage from a greater bid would be to BH with $v_{BH} = 3$ if could prevent being outbid by BL with $v_{BL} = 2$. If BH bid 2 in this instance, then it would make a guaranteed profit of 1, but by bidding $\frac{4}{3}$ it makes an expected profit of $\frac{5}{4} = \frac{5}{3}\frac{3}{4}$, due to being outbid 25% of the time. As a result, no first-round buyer benefits from bidding more than $\frac{4}{3}$. The strategies are therefore in PBE.

A simple way to look at the outcome of this equilibrium is to see that when $v_{BL} = 2$, the seller trades with whichever of BL or BH arrives first. When $v_{BL} = 1$, the low buyer's shading precludes the trade, and so the seller trades with the high buyer regardless of arrival order. Thus, $\frac{3}{4}$ of the time the seller sells to BH and the other $\frac{1}{4}$ the seller sells to BL when it has a valuation of 2. The expected welfare in this equilibrium is $2.375 = \frac{3}{4}2.5 + \frac{1}{4}2$. Strategic shading significantly improves the efficiency of this market from 0.86 (2.167 / 2.5) to 0.95 (2.375 / 2.5). However, equilibrium efficiency is not guaranteed to improve over truth-telling in CDAs, as demonstrated in my final stylized example.

Example 3. *Let there be two buyers and one seller, all with i.i.d. uniform private values in [0, 1]. The agents arrive in a uniform random ordering, and observe the current state of the order book, and the ordering of agents.*

Since the buyers are ex ante identical, I refer to the first buyer as buyer 1 (b_1) and the second buyer as buyer 2 (b_2).

The social optimum occurs when the seller trades with the largest buyer. If I let $\overline{v_b}$ represent the largest buyer valuation, then

$$\text{Optimal Welfare} = \mathbb{E}\left[\overline{v_b} - v_s \mid v_s < \overline{v_b}\right] = \frac{1}{4}.$$

If agents tell the truth, and the seller arrives last, it trades with the maximum of the buyer valuations, and so the welfare is the same as the optimal social welfare $\frac{1}{4}$. In the other two instances the seller has a chance to trade with the first buyer before trading with the second. Therefore

$$\begin{aligned}\text{Welfare}_{\textit{seller not last}} &= \mathbb{E}\left[v_{b_1} - v_s \mid v_s \leq v_{b_1}\right] + \mathbb{E}\left[v_{b_2} - v_s \mid v_s > v_{b_1}, v_s \leq v_{b_2}\right] \\ &= \frac{5}{24}.\end{aligned}$$

The efficiency of truth-telling is the expected welfare for truth-telling divided by the maximum social welfare, or $\frac{8}{9} = 4\left(\frac{2}{3}\frac{5}{24} + \frac{1}{3}\frac{1}{4}\right)$.

The PBE solutions can be calculated using backward induction. To calculate the efficiency in PBE, I first consider the case when the seller arrives first. In this case, both buyers will accept any offer below their valuation. The only strategic decision is the bid the seller should make. Using similar notation to the social optimum, the seller's expected profit for a bid s is

$$\text{Profit}_{\textit{seller}} = \mathbb{E}\left[s - v_s \mid s < \overline{v_b}\right] = (s - v_s)\left(1 - s^2\right).$$

This profit is maximized at $s^* = \frac{1}{3}\left(v_s + \sqrt{v_s^2 + 3}\right)$. The welfare calculation is the same as in the truthful case, except that the conditions are in terms of s^* instead of v_s. The social welfare when the seller arrives first is therefore $\sqrt{3}/9$.

When the seller arrives in the middle, the last buyer takes the sellers bid if it exists and the buyer can profit. The seller has a choice between taking the existing bid or placing a new one, attempting to get more surplus from the second bidder. The seller makes profit

$$\text{Profit}_s = \begin{cases} b_1 - v_s & \text{take order} \\ \mathbb{E}\left[s - v_s \mid s \leq v_{b_2}\right] & \text{place order } s. \end{cases}$$

The optimal bid in the later case is $s^* = \frac{1}{2}(1 + v_s)$, which implies that the seller takes the existing order if $v_s \leq 2\sqrt{b_1} - 1$, otherwise it places an ask at s^*. A placed ask will always be greater than the existing bid, that is, it never transacts with the old bid.

The first buyer's profit conditioned on this information is

$$\text{Profit}_{b_1} = \mathbb{E}[v_{b_1} - b_1 \mid v_s \le 2\sqrt{b_1} - 1] = (v_{b_1} - b_1)\,(2\sqrt{b_1} - 1).$$

Which is maximized when $b_1^* = \frac{1}{18}\,(1 + 6v_{b_1} + \sqrt{12v_{b_1} + 1})$, conditioned on $v_{b_1} \ge \frac{1}{4}$. The resulting social welfare conditioned of these strategies is $\frac{13}{5184}(47 + 8\sqrt{13})$.

When the seller arrives last, the first buyer knows it can win only by bidding over the second buyer's valuation, thus

$$\text{Profit}_{b_1} = \mathbb{E}[v_{b_1} - b_1 \mid v_{b_2} < b_1, v_s \le b_1] = b_1^2(v_{b_1} - b_1).$$

The optimal bid is $b_1^* = \frac{2}{3}v_{b_1}$. The second bidder will over bid the existing order if it can profit, but might bid more than epsilon over if it can extract more expected profit from the seller. The profit for placing a new order is

$$\text{Profit}_{b_2} = \mathbb{E}[v_{b_2} - b_2 \mid v_s < b_2, b_1 < b_2] = b_2(v_{b_2} - b_2).$$

This profit is maximized at $b_2^* = \frac{1}{2}v_{b_2}$ as long as $b_2^* \ge b_1$. Buyer 2's optimal strategy is to bid $\min\{v_{b_2}, \max\{b_2^*, b_1 + \epsilon\}\}$. The social welfare of this permutation is $\frac{19}{96}$, making the expected social welfare in PBE roughly 0.19.

The efficiency in PBE is the average of each permutation, which is approximately 0.77, significantly less than the corresponding truth-telling efficiency of approximately 0.89.

Of course, these are only particular instances; it is easy to construct other simple instances where strategic shading either helps or harms efficiency. In the remainder of this chapter I explore the effect of shading in richer scenarios, an expanded form of the model from Zhan and Friedman (2007) and a model more representative of trading situations that arise in financial markets. I overlay game-theoretic reasoning on a systematic simulation-based process, to investigate the impact of strategic bidding on outcomes realized in approximate equilibria, across a range of market environments.

4.3 Prior Work

There are many prior studies that attempt to characterize the efficiency of CDAs, however most do so from the perspective of semi-strategic agents, such as human lab participants or fixed algorithms designed without concern for equilibrium (Cason and Friedman 1996; Gjerstad and Dickhaut 1998; Rust, Miller, and Palmer 1993).

In contrast, Zhan and Friedman (2007) address the question with respect to Nash equilibrium shading. In their model, there are N buyers and N sellers with uniform i.i.d. private values for a single unit. The authors considered three different values of N: 4 (thin), 10 (medium), and 100 (thick); and three different classes of parameterized shading strategies: Standard, Exponential, and Shift. In order to analyze their model as a normal-form game, each strategy class was discretized into 11 strategies, with shading amounts uniformly spaced from 0 to 1. An agent's strategic choice is the amount of shading conditioned on a global shading class and market thickness. Agents arrive in a random ordering[1] and submit limit orders applying their shading strategies to their private values. In this model strategies do not consider market information, such as price quotes or transaction history; they are functions solely of private value.

Separate from their analysis of equilibrium shading, Zhan and Friedman used this model to explore how uniform non-strategic shading affects market efficiency. Their investigation of uniform shading in the thick-market standard-shading scenario follows the intuition from Example 1: that moderate shading mostly precludes inefficient trades. I replicated this experiment with my own implementation, and present the relationship between shading and efficiency in Figure 4.2 (which faithfully reproduces Zhan and Friedman's Figure 6).[2] Let IM inefficiency refer to the inefficiency caused because agents who would trade in competitive equilibrium do not (missing *intra-marginal* trades), and EM inefficiency refer to the inefficiency caused because agents who would not trade in competitive equilibrium do (present *extra-marginal* trades). The figure shows that as symmetric fixed shading increases from truth-telling, EM-inefficiency significantly decreases, while IM inefficiency remains close to zero. A little before 0.3 shading, both sources of inefficiency are minimized, and then as shading increases more, the IM inefficiency significantly increases. This suggests that the majority of inefficient trades have relatively small margins, and are inhibited by a small amount of shading, whereas the efficient trades have larger margins and so are uninhibited by modest shading.

In addition to analysis of the market environment under uniform shading, Zhan and Friedman also empirically found *role-symmetric pure-strategy* Nash equilibria in the nine variations of this game. In a role-symmetric strategy profile, each player within a role (here buyer or seller) plays the same pure or mixed strategy. Conveniently, each of their game instances had a single role-symmetric pure-strategy

[1] This description is slightly different from the original, but produces identical results.

[2] This is true despite the fact that I use a different definition of expected efficiency, stated in the beginning of Section 4.2. Section 4.4 discusses this difference in more detail.

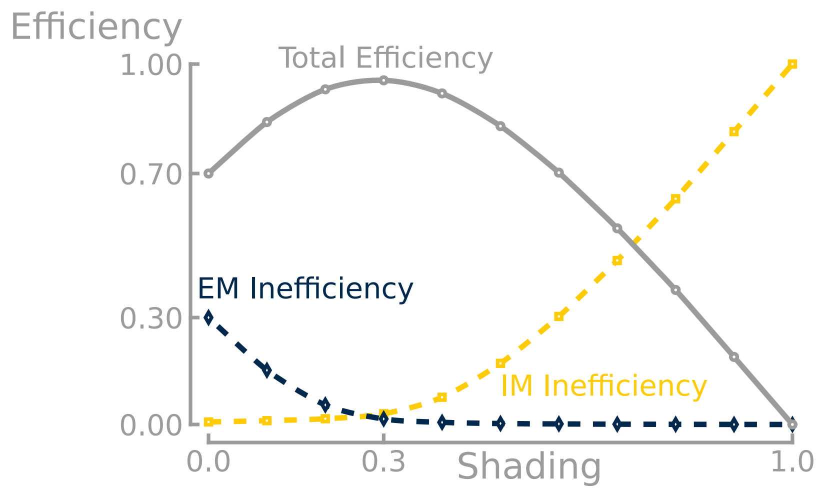

Figure 4.2: Moderate shading reduces inefficiency from suboptimal trades (EM-inefficiency) without introducing significant inefficiency from missed optimal trades (IM inefficiency). The data presented is from a replication of Zhan and Friedman (2007, Figure 6).

equilibrium, or a single ε-approximate one for small ε. Zhan and Friedman conclude from the efficiency of these equilibria that CDA equilibrium surplus is *close* to optimal surplus. However, their results for the thin market—with four buyers and four sellers—indicate that the equilibrium surplus, while high, can be much worse than truth-telling. The issue of non-universal improvement over truth-telling is not discussed much by Zhan and Friedman, but is the focus of my extended investigation. In Section 4.4 I present the rest of the results of my replication and extension of this work.

4.4 Replication of Zhan and Friedman

I first present a replication of Zhan and Friedman (2007), which extends their study in three key ways:

1. The original study separately analyzed games with three classes of discretized strategies. I investigated a fourth strategy class where agents can choose any strategy from the original three classes.

2. The original study restricted solutions to ε pure-strategy Nash equilibria. I broadened consideration to include mixed-strategy equilibria, and accordingly searched more extensively over strategy profiles.

3. I evaluated the found equilibria with many more samples and report statistical

confidence on regret.

My replicated CDA market simulator is identical to that specified by Zhan and Friedman except for one detail. Instead of allowing agents to rebid at the same price, I simply shuffle the agents once and have them submit bids in that order. Aside from subtle effects on time priority, this should produce identical results. I confirm empirically that any differences are negligible.

Also, while the precise aggregate efficiency measure used was unspecified in the original paper, I were able to reproduce the reported efficiency results exactly only by calculating averages over instance efficiencies (i.e., surplus obtained as a fraction of CE surplus for each instance). I argue that the proper measure of expected efficiency is the expected surplus over the expected CE surplus. This definition appropriately gives more weight to random instances that allow more surplus and removes the necessity of defining efficiency in a no-trade scenario.

With this simulator, I applied the equilibrium search methodology presented in Chapter III using the same strategy discretization as the original authors. I evaluated at a total of twelve scenarios formed by three levels of market thickness combined with four classes of shading strategies. Three of the strategy classes—Standard, Exponential, and Shift—were introduced by Zhan and Friedman. I added a union shading class ("All"), containing 31 strategies: eleven shading levels from each original class, minus two that are redundant because zero shading corresponds to truth-telling in every shading class. In order to tractably explore each game, I applied DPR, and reduced the number of players in each role to four.[3] To consider a profile in each of these games explored, I sampled it 250,000, 50,000, and 25,000 times respectively for the thin, medium and thick markets, or an order of magnitude more samples than the previous study. Finally, I set the restricted game size limit for equilibrium search to four, meaning I stopped exploring profiles after I had found at least one equilibrium and all unexplored best response restricted games had at least five strategies in support between both roles.

A summary of the experimental results is presented in Figure 4.3. The equilibrium efficiency between shading classes in a single market thickness varies slightly. This oscillation is likely due to a few factors including the choice of discrete strategies,[4] the sampling error inherent in random simulation, and the fact that my profile search biases towards low support equilibria. Despite these factors, there is a strong overall

[3]The thin market games were unreduced.

[4]In particular, any Exponential shading can be achieved by a Standard shading at some transformed level, so the essential difference between the classes is really the choice of discrete shading levels.

trend in the data. As the market gets thinner, the efficiency of truth telling improves, and that of equilibrium shading degrades.

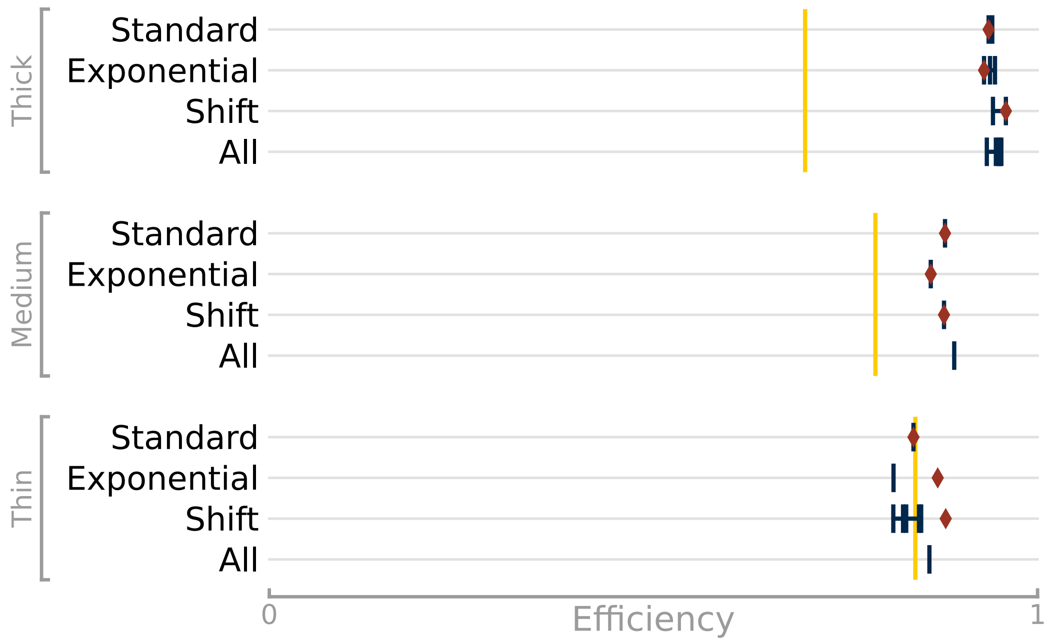

Figure 4.3: Efficiencies of all found equilibria compared to truth-telling in every scenario. Yellow bars represent the truth-telling efficiency in each market thickness. Blue bars with a line connecting them represent the efficiency of every role-symmetric equilibria found via my profile exploration method. Red diamonds represent the efficiencies of each symmetric pure-strategy equilibrium found in Zhan and Friedman (2007). Efficiency is computed as the mean surplus over the mean competitive equilibrium surplus for a large sample of profiles drawn from the equilibrium distribution.

In addition, my search process failed to find two equilibria from the previous work. However, not all found equilibria are equally important. Approximate equilibria with low regret are more plausible descriptions of rational agent behavior than those with high regret, and while I cannot confirm that regret of any equilibrium is low, I can compute confidence bounds on the regret. Table 4.1 lists all of the equilibria, their efficiency, and a 95% bootstrap upper confidence bound on regret. The previous equilibria that I did not find have significantly higher—roughly double—the regret bound of my comparable found equilibria.

4.5 Financial Market Environment

My extension enriches the CDA scenario to more closely resemble current financial markets by using a modification of the model presented in Chapter II. Agents observe

Table 4.1: Equilibria found in CDA games, by Zhan and Friedman (2007) or my replication. Where multiple equilibria were found, they are given numbers to differentiate. Equilibria numbered "ZF" correspond to equilibria found in the previous work, but not this one, while equilibria labeled with a † were also identified in the previous work. In the "All" shading class agents can play any strategy from Standard, Exponential, or Shift. To differentiate these strategies in the "All" class, standard shading strategies have no suffix, exponential strategies have an "E" suffix, and shift strategies have an "S" suffix. "95% Regret" is the bootstrapped 95% upper confidence interval on regret.

Density	Shading Class	Num	Buyers' Strategies	Buyers' Prob (%)	Sellers' Strategies	Sellers' Prob (%)	Efficiency	95% Regret
Thick	Truthful	–	0.0	100	0.0	100	0.697	–
	Standard	1†	0.4	100	0.3	100	0.936	0.466
		2	0.4	100	0.4, 0.2	69.4, 30.6	0.941	0.235
	Exponential	1†	0.5	100	0.3	100	0.930	0.521
		2	0.5, 0.6	73.9, 26.1	0.2	100	0.945	0.198
		3	0.5	100	0.2, 0.3	38.7, 61.3	0.938	0.132
	Shift	1†	0.4	100	0.4	100	0.959	0.255
		2	0.4	100	0.3, 0.5	12.5, 87.5	0.942	0.537
	All	1	0.4	100	0.2S	100	0.951	0.845
		2	0.4E	100	0.4S, 0.3S	86.9, 13.1	0.953	0.566
		3	0.3, 0.4, 0.5E	16.1, 27.9, 56	0.3S	100	0.934	0.910
		4	0.4, 0.4E, 0.5E	12.8, 63, 24.2	0.3S	100	0.946	0.421
		5	0.4	100	0.3S, 0.2S	95.9, 4.1	0.950	0.423
Medium	Truthful	–	0.0	100	0.0	100	0.789	–
	Standard	1†	0.3	100	0.3	100	0.880	0.045
	Exponential	1†	0.4	100	0.3	100	0.861	0.030
	Shift	1†	0.4	100	0.4	100	0.878	0.000
	All	1	0.3	100	0.3S	100	0.892	0.165
Thin	Truthful	–	0.0	100	0.0	100	0.841	–
	Standard	1†	0.3	100	0.3	100	0.838	0.151
	Exponential	ZF	0.3	100	0.2	100	0.870	0.565
		1	0.4	100	0.3	100	0.812	0.066
	Shift	ZF	0.3	100	0.3	100	0.880	0.285
		1	0.3	100	0.4	100	0.848	0.043
		2	0.4	100	0.4, 0.3	93.1, 6.9	0.846	0.198
		3	0.3, 0.4	43.7, 56.3	0.4	100	0.825	0.098
		4	0.3, 0.4	9.8, 90.2	0.4, 0.3	2.3, 97.7	0.812	0.180
		5	0.3, 0.4	11.8, 88.2	0.4, 0.3	42.5, 57.5	0.829	0.094
	All	1	0.3E	100	0.3S	100	0.859	0.032

a positive integer truncation of the fundamental $\tilde{f}_t = \left\lfloor f_t + \frac{1}{2} \right\rfloor_+$, which is nearly identical given the high shock variance and large mean in the analyzed environment settings. In addition, I consider two cases for order submission. In the first (*single-unit order*) case, agents are assigned on each arrival to either buy or sell, with equal probability, and accordingly submit an order to buy or sell a single unit. This style of order submission was described in base market model in Section 2.2. In the second (*full-demand order*), agents submit orders to both buy and sell. Agents with this style of order submission are not subject to the random direction selection of the single-unit order case.

Since agents withdraw stale orders at arrival, the arrival rate λ serves as a rough proxy for agent urgency. If an agent has a lower arrival rate, then it has fewer opportunities to submit orders, which is particularly constraining in the single-unit order case. With single-unit ordering, agents may find it difficult or impossible to achieve their efficient position levels. In such settings, they have a strong incentive to make each arrival count. In the full-demand order case agents can achieve their efficient position levels in one arrival, thus providing a way to evaluate the effect of reducing urgency.

Similar to Zhan and Friedman (2007), I restrict my analysis to a discretized finite set of the entire strategy space—shown in Table 4.2—so I can apply normal-form game analysis techniques. I selected strategies using the technique detailed in Section 3.1, with the hope that equilibria in the normal-form game exhibit low regret in the continuous game. The initial set of strategies was chosen as a roughly exponential grid of the max shading ($\overline{d}$) parameter (Table 4.2a). Strategies that appeared frequently as best responses to equilibria of games with this initial set were also considered (Table 4.2b). The union of these two sets has thirteen strategies, which are the only strategies I consider for the financial market games. Discretization provides no guarantees about regret in the continuous game, but it does allow me to make precise statements about regret in this specific normal-form game.

Generalizing the ZI strategy to multiple units in the same direction is not obvious, so I consider only a restrictive version of the full-demand order case. Specifically, my full-demand environments assume a maximum absolute position ($q_{\max}$) of one, in which case when agents hold no position, they submit a full demand schedule. Agents need to arrive only once to reach CE position. The agent draws its desired surplus (shading) once per arrival, and applies it to both trade directions.

Table 4.2: Strategies considered for equilibrium analysis. Agents shade from their true belief by a random offset in $[\underline{d}, \overline{d}]$, the result of this draw is the amount of expected surplus the agent demands from trade. If an agent could get an η fraction of their desired surplus from an outstanding order, they will take the order instead. Detailed descriptions of the strategies can be found in Section 2.2.

(a) Initial set of strategies at roughly exponentially spaced intervals.

$\underline{d}$	$\overline{d}$	η
0	0	1
0	50	1
0	125	1
0	250	1
0	500	1
0	1000	1
0	2000	1

(b) Best response strategies, found via iterative process of adding best response to previous equilibria.

$\underline{d}$	$\overline{d}$	η
125	500	1
2000	4000	1
0	500	0.8
0	2000	0.8
125	1000	0.4
29000	30000	0.001

4.5.1 Market Environment Parameter Settings

I focus my analysis on two basic market settings, with parametric variations. Restricting my analysis to these two settings allows me to investigate qualitative effects without having to do an infeasible exhaustive grid search over the space of parameter configurations. Baseline parameter values for these settings are available for reference in Table 4.3. In both settings I vary the arrival rate, mean reversion, and market thickness.

The *multi-position* setting most closely matches the environment for which ZI agents were originally defined (Gode and Sunder 1993). Agents have a maximum absolute position of ten (hence the name multi-position), and follow the single-unit order scheme. As a result, the agents never have more than one outstanding single-unit order in the market at a time.

The second setting I call the *single-position* setting. In this setting, agents have a maximum position of one, and follow the full-demand order scheme. This setting lifts the restriction that agents have exactly one outstanding single-unit order at a time, but agents will never hold an absolute position greater than a single unit (hence the name single-position).

4.5.2 Uniform Shading Analysis

One would hope that the intuition behind the tradeoff of EM and IM inefficiency in the Zhan and Friedman model—seen in Figure 4.2—would carry over to this financial

Table 4.3: Baseline environment parameters, by setting. Arrival rate is double in the multi-position setting because agents need to trade more in competitive equilibrium.

Parameter	Symbol	Multi-position	Single-position
Number of Agents		66	
Trading Horizon	T	60000	
Fundamental Mean	$\bar{f}$	10^7	
Mean Reversion	r	0.05	
Shock Variance	σ_s^2	10^6	
Private Value Var	σ_{pv}^2	5×10^6	
Arrival Rate	λ	10^{-4}	5×10^{-5}
Maximum Position	$q_{\max}$	10	1
Order Submission		Single-unit	Full-demand

market model as well. Figure 4.4a is a plot of EM and IM inefficiency for the baseline multi-position setting. Unlike the standard CDA model, EM and IM inefficiency do not provide clear clues as to *how* shading affects efficiency in this model. I hypothesize that the definitions of EM and IM inefficiency are muddled when agents have the opportunity to trade counter to their competitive equilibrium, that is, an agent who *should* buy a unit might opportunistically sell one due to advantageous market conditions. In this environment, EM trades happen when an agent trades beyond their CE position—or at all if their CE position is zero. When every agent submits truthful orders, the incentive to take an opportunistic trade overshadows the incentive to trade towards competitive equilibrium, resulting in agents not achieving their desired positions. The incentive to trade counter to competitive equilibrium produces high IM inefficiency even when agents are truthful, but moderate shading reduces this effect slightly. Figure 4.4b is an identical plot of EM and IM efficiency for the baseline multi-position setting, except half of the agents were assigned to be buyers, half were assigned to be sellers, and each agent could only trade in their assigned directions. When this structure is applied to the market, the illustrative breakdown of efficiency presented in the standard CDA model (Figure 4.2) is recovered. It's clear from these figures that the ability for agents to trade in both directions adds a significant level of complexity to analyzing the effect of shading on efficiency in CDAs. I leave a more appropriate dissection of efficiency when agents can trade in both directions as an open problem, and instead focus on the effects of shading in equilibrium.

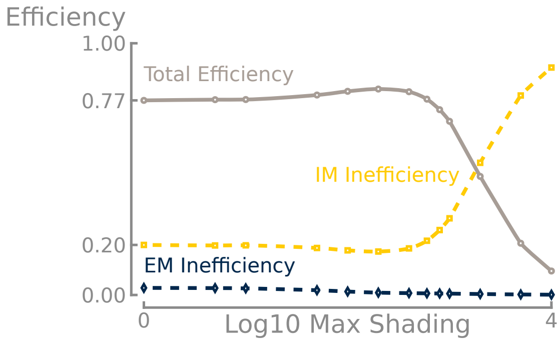

(a) EM and IM inefficiency in the financial market model. Unlike the standard CDA model, this decomposition does not present a clear picture of *how* shading affects efficiency.

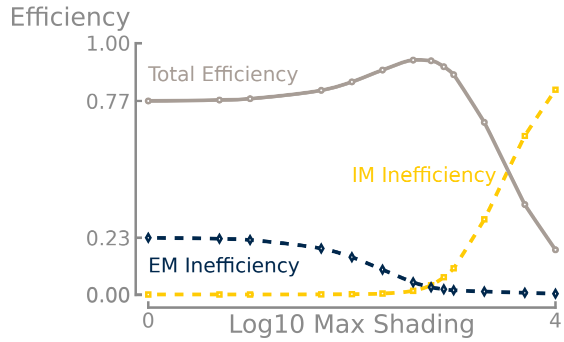

(b) EM and IM inefficiency in the financial market model when agents are assigned to be buyers or sellers. With this condition, the effect of shading on efficiency is qualitatively similar to the results from Zhan and Friedman (2007) (Figure 4.2).

Figure 4.4: Uniform shading analysis of the financial market model

4.6 Financial Market Equilibrium Analysis

Using the model described in Section 5.3, and the methodology to find approximate role-symmetric Nash equilibria described in Chapter III, I investigate the effects of equilibrium shading and truth-telling on financial market efficiency. To measure efficiency in this model, I calculate an agent's competitive equilibrium position constrained by its number of arrivals. This is necessary since some agents may not arrive at all. Due to the large number of agents in the simulation, I used DPR to reduce the effective number of agents to six. I considered a profile explored if I sampled it 10,000 times, at this level, the standard error of profile payoffs was sufficiently small. I set the restricted game size limit to three, meaning I stopped exploring profiles after I had found an equilibrium, and all best response restricted games had support over at least four strategies.

In both the single- and the multi-position setting I varied three key parameters from their baseline value: number of agents, arrival rate, and fundamental mean reversion. Increasing the number of agents in a simulation increases the market thickness. Increasing the arrival rate increases agents' access to relevant information, decreases the amount of time potentially stale order sit in the market, and gives agents more opportunities to trade. Finally, increasing the mean reversion exposes agents to less adverse selection due to the smaller impact of a shock on the final fundamental price. In each setting, I analyze twelve distinct environments corresponding to different levels of these three parameters. With no mean reversion and other settings at baseline, the adverse selection is severe enough to preclude all trading. To restore some profitability in this environment, I compensate by also reducing the shock variance σ_s^2 to 100.

Figure 4.5 compares the efficiency achieved in equilibrium with that produced by truth-telling agents for each parameter variation. Both the single- and multi-position setting show identical trends for each variable. Increasing the number of agents increases the benefit from equilibrium shading. This is the same trend that was found in the simple CDA model. Unlike that model, truth-telling efficiency also increases with the number of agents. This is probably because agents in this model can retrade, effectively making initial extra-marginal trades less important to final efficiency. Increasing either arrival rate or mean reversion has the opposite trend, the benefit to equilibrium shading decreases, despite an increase in truth-telling efficiency. This is to be expected, as a limited number of trader arrivals means little time to correct for bad trades and the increase in adverse selection from low mean reversion tends to generate more inefficient trades. Shading ameliorates inefficient trades, yielding an improvement in efficiency in both environments.

Tables 4.4 and 4.5 list all of the equilibria I found. None of the equilibria had large regret upper confidence bounds, so while DPR may have affected which equilibria I found, the statistical evidence suggests that it did not degrade the *quality*—low regret—of the equilibria I found. In addition, in games where I found both pure and mixed-strategy equilibria, the upper confidence regret bound on mixed-strategy equilibria was always close to the confidence bound for pure-strategy equilibria. Since pure strategy equilibria found by DPR are also equilibria in the unreduced game, this evidence suggests that the regret caused by DPR may be small compared to the inherent variability of the models.

4.7 Conclusions

I employed a simulation-based approach to analyze the effect of strategic bid shading on the efficiency of standard CDA markets and richer financial markets. I confirmed most of the results from Zhan and Friedman (2007) using a more complete equilibrium search. In both market models that I investigated, I consistently found that strategic bid shading helps efficiency when there are more agents in the market. In the financial market, I also observe a benefit to strategic shading when there is a large amount of adverse selection due to a noisy common valuation or a high level of urgency due to limited trading opportunities. My results strengthen the claims made

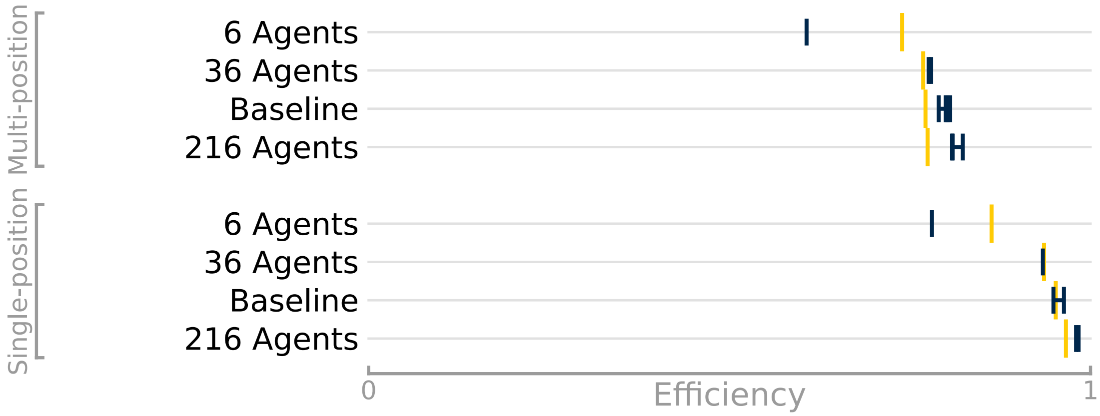

(a) Equilibrium efficiency as a function of number of agents.

(b) Equilibrium efficiency as a function of arrival rate.

(c) Equilibrium efficiency as a function of mean reversion.

Figure 4.5: Fraction of optimal surplus for truth-telling agents (yellow), and the efficiencies in equilibrium (blue) for each of the environments. The equilibrium values are displayed as a range because multiple equilibria were found in each environment.

Table 4.4: Equilibria found in multi-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered "T" correspond truth-telling profiles. "95% Regret" is the bootstrapped 95% upper confidence interval on regret.

Environment	Num	$U_{\min}$	$U_{\max}$	η	Prob (%)	Efficiency	95% Regret
4/5×Arrival Rate	T	0	0	1	100	0.702	–
	1	0	500	0.8	100	0.743	33.954
	2	0,125	500,500	0.8,1	19.4,80.6	0.734	23.894
6 Agents	T	0	0	1	100	0.739	–
	1	0,125,125	500,500,1000	1,1,0.4	2.2,3.5,94.3	0.607	3.486
36 Agents	T	0	0	1	100	0.768	–
	1	0,0,125	500,500,500	0.8,1,1	25.3,10.4,64.3	0.776	24.880
	2	0,125	500,500	1,1	67.5,32.5	0.779	19.546
0×Mean Reversion	T	0	0	1	100	0.469	–
	1	125,29000	1000,30000	0.4,0.001	55.2,44.8	0.602	33.773
1/10×Mean Reversion	T	0	0	1	100	0.757	–
	1	0,125,125	500,500,1000	1,1,0.4	3.5,46.6,49.9	0.761	47.824
1/5×Mean Reversion	T	0	0	1	100	0.766	–
	1	0,125	500,500	0.8,1	2.3,97.7	0.783	54.921
	2	0,0,125	250,500,1000	1,1,0.4	26.3,63,10.6	0.793	73.490
2/5×Mean Reversion	T	0	0	1	100	0.770	–
	1	0,0	250,500	1,1	70.9,29.1	0.811	26.887
	2	0,125	500,1000	0.8,0.4	97.5,2.5	0.797	63.857
	3	0,125,125	500,500,1000	0.8,1,0.4	3.1,82.1,14.8	0.781	22.112
	4	0,0,125	125,500,500	1,0.8,1	10.7,45.1,44.1	0.795	21.439
Baseline	T	0	0	1	100	0.771	–
	1	0,0	250,500	1,0.8	19.3,80.7	0.803	16.131
	2	0,0,125	250,500,500	1,0.8,1	8.5,90.6,0.9	0.800	15.544
	3	0,0,125	250,500,500	1,0.8,1	29.6,27.1,43.3	0.800	17.152
	4	0,0,0	125,250,500	1,1,1	20.8,23.4,55.7	0.805	17.261
	5	0,0,125	500,500,500	0.8,1,1	21.4,10.9,67.7	0.790	3.687
2×Arrival Rate	T	0	0	1	100	0.955	–
	1	0,125	250,1000	1,0.4	94.7,5.3	0.957	24.640
	2	0,29000	500,30000	1,0.001	87.3,12.7	0.941	43.369
	3	0,125	500,500	1,1	24.1,75.9	0.923	16.494
	4	0,0,0	50,125,500	1,1,0.8	8.7,7.7,83.6	0.945	38.080
	5	0,0,125	500,500,500	0.8,1,1	36.7,13.1,50.1	0.931	26.133
5×Arrival Rate	T	0	0	1	100	0.997	–
	1	0,0	50,500	1,0.8	9.5,90.5	0.990	40.257
	2	0,125	1000,500	1,1	7.1,92.9	0.963	62.199
10×Arrival Rate	T	0	0	1	100	0.994	–
	1	0,125	1000,500	1,1	7.9,92.1	0.970	51.134
	2	0,0,125	250,2000,500	1,0.8,1	0.5,7.6,91.9	0.969	25.813
216 Agents	T	0	0	1	100	0.774	–
	1	0,0	250,500	1,1	78.1,21.9	0.823	19.123
	2	0,125	500,1000	0.8,0.4	98.1,1.9	0.809	41.196
	3	0,0,125	500,500,500	0.8,1,1	45.2,23.8,31	0.808	43.869

Table 4.5: Equilibria found in single-position financial market environments. Where multiple equilibria were found, they are given numbers to differentiate. Profiles numbered "T" correspond truth-telling profiles. "95% Regret" is the bootstrapped 95% upper confidence interval on regret.

Environment	Num	$U_{\min}$	$U_{\max}$	η	Prob (%)	Efficiency	95% Regret
$4/5\times$Arrival Rate	T	0	0	1	100	0.938	–
	1	125	500	1	100	0.958	10.511
	2	0,125	1000,1000	1,0.4	31.7,68.3	0.947	8.514
6 Agents	T	0	0	1	100	0.863	–
	1	0,125,29000	2000,1000,30000	0.8,0.4,0.001	49.1,45.7,5.2	0.780	8.143
36 Agents	T	0	0	1	100	0.936	–
	1	125,125	500,1000	1,0.4	9.1,90.9	0.934	0.958
$0\times$Mean Reversion	T	0	0	1	100	0.736	–
	1	0,29000	2000,30000	0.8,0.001	66.6,33.4	0.824	35.347
$1/10\times$Mean Reversion	T	0	0	1	100	0.925	–
	1	0	2000	0.8	100	0.880	45.380
	2	0,125	2000,1000	0.8,0.4	84.9,15.1	0.889	29.130
	3	0,125	2000,1000	0.8,0.4	16.8,83.2	0.926	29.222
$1/5\times$Mean Reversion	T	0	0	1	100	0.942	–
	1	0,125	1000,1000	1,0.4	22.9,77.1	0.947	11.404
	2	0,125	500,500	0.8,1	28.2,71.8	0.956	29.546
$2/5\times$Mean Reversion	T	0	0	1	100	0.949	–
	1	0,125	1000,1000	1,0.4	44.9,55.1	0.951	2.847
	2	0,0,125	500,500,500	0.8,1,1	33.7,4.7,61.5	0.961	17.122
Baseline	T	0	0	1	100	0.952	–
	1	0	1000	1	100	0.949	14.405
	2	0,125	500,500	1,1	22.7,77.3	0.963	13.473
$2\times$Arrival Rate	T	0	0	1	100	0.976	–
	1	125,125	500,1000	1,0.4	58.2,41.8	0.969	7.768
$5\times$Arrival Rate	T	0	0	1	100	0.975	–
	1	0,0,125	1000,2000,500	1,1,1	0.2,90.8,9	0.964	15.971
	2	0,125	1000,1000	1,0.4	34.4,65.6	0.972	6.737
	3	0,125,125	1000,500,1000	1,1,0.4	17.7,22.3,60	0.972	11.430
	4	0,125,125	2000,500,1000	0.8,1,0.4	15.4,13.2,71.3	0.970	11.174
$10\times$Arrival Rate	T	0	0	1	100	0.971	–
	1	0,0	1000,2000	1,1	68.4,31.6	0.971	19.954
	2	0,125	2000,1000	1,0.4	24.7,75.3	0.970	15.222
	3	0,125	2000,1000	0.8,0.4	22.1,77.9	0.970	13.117
216 Agents	T	0	0	1	100	0.966	–
	1	125	500	1	100	0.983	7.130
	2	0,125,125	500,500,1000	1,1,0.4	89.5,0.7,9.8	0.980	14.469

by Zhan and Friedman by broadening the search for Nash equilibria and extending the model to environments with dynamic elements.

More generally, this phenomenon is highly germane to design of rules and regulations surrounding CDA markets. For example, measures aimed at promoting true value revelation—following the typical intuition of mechanism designers—may be counter-productive to mechanism design goals in this setting. My analysis also underscores the need for accounting for strategic behavior when comparing CDAs to alternative mechanisms, particularly those (e.g., call markets) that may not exhibit this phenomenon.

CHAPTER V

Empirical Mechanism Design for Optimizing Clearing Interval in Frequent Call Markets

5.1 Introduction

Most modern financial markets operate a continuous double auction (CDA). A consequence of the continuous-time property of CDAs is that there is no lower bound on the speed differential that could be pivotal in deciding the trade. Financial markets adopting this mechanism foster a *latency arms race*, where traders may obtain significant benefit by achieving small access-time advantages over their counterparts.

The apparent first-order cost of the latency arms race has led several to advocate switching from continuous to discrete-time clearing mechanisms (Budish, Cramton, and Shim 2015; Schwartz and Wu 2013; Sparrow 2012; Wah and Wellman 2013). In a discrete-time, or *call market*, orders are accumulated (or *batched*) over specified time intervals, with matches determined only at the end of each interval.

The call market mechanism has been known and in substantial use for a long time; in fact typical financial markets open each day with a call market. What is new is the opportunity to run call markets with short clearing intervals, say on the order of one second. The Chicago Stock Exchange recently launched a similar mechanism called a SNAP^{SM} auction,[1] which is roughly a temporary half-second call that occurs in an otherwise standard CDA market. The clearing interval in frequent call markets is intended to be long enough to undermine most of the benefit of tiny speed advantages, yet short enough to avoid imposing serious delays on economic activity. The literature discussing discrete-time markets uses a number of different names, including frequent batch auction (Budish, Cramton, and Shim 2015), periodic auction (Madhavan 1992),

[1] `http://www.chx.com/snap/snap-auction-cycle.html`

and frequent call market (Wah and Wellman 2013); I simply refer to them as call markets, with the understanding that the clearing interval will be relatively short.

There are many ways a call market could trigger a clear, such as by volume or time, deterministically or randomly. Random choice of clear time has been proposed as a way to prevent a form of *sniping*, where a fast trader with an edge on information or market access can snipe a stale order. In a call market, this applies only right at the end of the clearing interval, say for a fraction δ of that interval. Randomization prevents such deterministic sniping, but allows the fast trader to snipe with a δ probability of success for the entire time. Thus randomization provides no protection benefit in expectation. It also poses the problem of trusting the mechanism to honestly randomize, which is not easy to audit.

For both of these reasons I focus on deterministic clearing, and the question of how long the interval should be. A specific choice of around one second, which has often been suggested as a reasonable interval, is quite arbitrary. Market designers should prefer a principled basis for deciding exactly how much time should elapse between market clears. I investigate this question by examining how the quality of the market—measured as allocative market efficiency and price discovery—is affected by the clearing interval when traders are strategic and utility maximizing, playing according to an approximate form of Nash equilibrium.

My model posits finitely many agents trading a single security in a single market, where agents respond to asymmetric private information as they repeatedly enter the market to buy or sell, conditioning on market observations. This model allows me to investigate how the presence of adverse selection, and market parameters that affect it, contribute to the optimal selection of a clearing interval. I do not model agents' response to public information or other important aspects of modern financial markets such as opportunities to invest in better information or faster market access. We should expect the duration of the clearing interval to have certain qualitative effects on a call market. If the clearing interval is sufficiently short, then a frequent call market should behave almost identically to a CDA.[2] Frequent clearing reduces transaction times at the expense of making hasty matching decisions and thereby potentially inefficient trades. If the clearing interval is sufficiently long, then the behavior approaches a *one-shot double auction* (Satterthwaite and Williams 1989; Wellman 2011). Infrequent clearing produces more efficient trades at the expense of

[2]Even at infinite frequency, the behavior is not exactly identical, because the markets price differently. CDAs price at the incumbent order price, whereas call markets price at the midpoint of the clearing range.

delayed execution. Waiting for the next clear can be costly, for example when other market actions serially depend on the present one. At some point, the efficiency benefit of longer clears will be outweighed by this time cost. The question then becomes, is there some clearing interval where the benefit of more efficient trades outweighs the loss of immediacy, or is the most efficient call market the one that behaves like a CDA?

My market model is too complex to analyze using traditional means, so I use the techniques from Chapter III to find approximate equilibria in a CDA market and in call markets with a discrete set of clearing intervals. In effect, I perform *empirical mechanism design* (Vorobeychik, Kiekintveld, and Wellman 2006) with the goal of maximizing allocative efficiency on a much richer model than would be possible to analyze with standard economic analysis. My analysis indicates that frequent batch auctions do not always improve market efficiency, but in markets with low adverse selection, lengthening the clearing interval produces greater efficiency until a point where agent impatience dominates and efficiency falls sharply. I also find that the benefit of switching to a frequent call market is significantly improved in thick markets. These conclusions suggest practical considerations that should be taken when implementing such a mechanism.

5.2 Related Work

Wah and Wellman (2013) and Budish, Cramton, and Shim (2015) were among the first to propose call markets as a design response to the negative effects of latency arbitrage. Both papers demonstrate that continuous clearing creates arbitrage opportunities that arise at imperceptible frequencies, and that exploiting them reduces market liquidity and efficiency. Budish, Cramton, and Shim (2015) also roughly calculate what batch interval should be sufficient to prevent latency arms races for various calibrations of their model. Their results yield a lower bound of one millisecond to one second depending on assumptions about the information arrival process and the magnitude of speed improvement available to high-frequency traders (HFTs). Other studies have also found benefits to frequent call markets. Schnitzlein (1996) performed a laboratory experiment comparing a call markets and CDAs and found that the call market was more liquid and better for noise traders, while not harming price efficiency. Baldauf and Mollner (2015) compared CDAs, frequent call markets, and delayed CDAs[3] in a model with costly research. Both frequent call markets and

[3]CDAs with delayed submission but immediate cancels.

delayed CDAs performed at the Pareto frontier of price discovery and spread, but traditional CDAs did not. Madhavan (1992) considered a model where risk-averse traders start with some initial endowment of both a risky asset and cash and found that, in equilibrium, call markets offered greater price efficiency than CDAs. Follow-up studies have investigated how frequent call markets could coexist with traditional CDAs (Li and S. Das 2016; Wah, Hurd, and Wellman 2015), have outlined practical considerations necessary to implement a frequent call market compliant with current regulations (Budish, Cramton, and Shim 2014), and present some tangential evidence of a smooth tradeoff between immediacy and batching (Wah and Wellman 2016). All these have relegated the question of optimizing call frequency to future work.

In response, several recent papers have developed stylized analytical models that propose definitions of an optimal clearing interval. Fricke and Gerig (2018) studied a financial market model that is very similar to the model used in Garbade and Silber (1979). Like the earlier paper, the authors attempted to determine the optimal clearing frequency based on a version of *liquidity risk*. The authors then calibrated their model with aggregate current market data and estimated that the optimal clearing interval for current S&P 500 securities is somewhere between one and three seconds. Du and Zhu (2017) solve for Nash equilibrium in linear strategies of a model with a divisible good and shocks to both common and private information. In this model, faster auctions allow agents to more quickly respond to information, but also cause them to be more aggressive with their orders. A key result is that when information arrives via a Poisson process, then the optimal clearing rate can be much slower than the average arrival rate of information, except when there are few traders or a large amount of adverse selection. Their approach is similar to mine, but several aspects of the model, including divisible units, private value shocks, and maintained order books, make it difficult to directly compare results. Haas and Zoican (2016) propose a model in which they measure market liquidity as a function of clearing interval in equilibrium. Their market model has two types of participants: HFTs and liquidity agents, where HFTs respond to random information and the liquidity traders are impatient and respond to random private value shocks. Their model parallels Budish, Cramton, and Shim (2015) until the clearing interval gets long enough, at which point traders' impatience gives out and longer clearing intervals cease to improve liquidity. The authors find that liquidity may still be worse in these call markets than in traditional CDAs.

5.3 Market Model

My experiments employ a parameterized financial market model, extended from the base model described in Chapter II by adding a call market. The market reveals price quotes reflecting the best outstanding orders after the most recent clear; other bids in the market are not visible to agents, and *no* new information is provided between clears. During a market clear, the mechanism determines the range of prices that would match supply and demand, and executes all compatible trades at the midpoint of this range. If there is an imbalance of supply and demand at the clearing price (due to multiple units with the same price), the ties are broken first by orders that arrived in earlier clearing intervals, and then in a random manner. CDAs are a special case of call markets that clear after every order submission and price transactions at the incumbent order price instead of the midpoint.

5.3.1 Call Market Trading Strategies

Since the ZI family of trading strategies was conceived for CDA markets, it stands to reason that effective call market strategies would have a different parameterization. To my knowledge, there is no prior literature analyzing strategies for periodic call markets. I propose a simple extension to the ZI class of strategies that allows agents to condition their shading on their location in the clearing interval. My extension provides the first step in exploring effective call market strategies. An agent that arrives close to a clear has better information about the fundamental than the other agents with orders in the market, and as a result, faces lower adverse-selection. Instead of being parameterized by only $\underline{d}$ and $\overline{d}$, this strategy class has four shading parameters $\underline{d}_i$, $\underline{d}_f$, $\overline{d}_i$, and $\overline{d}_f$, where i and f stand for *initial* and *final* respectively. When an agent arrives, it linearly interpolates its effective $\underline{d}$ and $\overline{d}$ based on its position in the interval. That is, if an agent arrives α fraction of the way through the interval, then its demanded surplus is drawn uniformly from the range $[(1-\alpha)\underline{d}_i+\alpha\underline{d}_f, (1-\alpha)\overline{d}_i+\alpha\overline{d}_f]$.

5.3.2 Environments

Characterizing all possible model parameter settings under all possible agent strategies is infeasible and would be largely uninteresting. Many possible parameter settings are degenerate—they provide no opportunities to trade, or insufficient information to trade safely, or fail to present strategic tradeoffs. For example, as fundamental shock variance diminishes, adverse selection vanishes and almost any trading strategy is sufficient to reach an efficient allocation, but if fundamental shock variance is too high, the shading required to leave a profitable resting order will make transacting unlikely. My analysis provides an understanding of how different aspects of a market affect optimal mechanism design by analyzing many discrete parameter settings. I refer to each distinct point in parameter space as an *environment*.

Each environment I study shares a base set of parameter settings. These fixed parameters either have an unimportant effect on the simulation, or can be appropriately scaled by settings of the parameters I do vary. The fundamental mean $\bar{f}$ is set very high at 10^9 so that boundary effects at zero never happen in practice. The private value variance σ_{pv}^2 is 5×10^6 and the maximum position $q_{\max}$ is 10 so that the importance of public and private value components are loosely balanced. With $q_{\max} = 10$, there is a 0.0002% chance that an agent's optimal position is 10, thus agents rarely reach their maximum position. The arrival rate λ is 5×10^{-3}. The remaining parameters are dimensions that I vary in my analysis: the number of agents (market thickness), the simulation length (number of trade opportunities), and the fundamental parameters (mean reversion, shock variance). I explore three different numbers of agents: 25, 58, and 238 respectively referred to as the thin ($\downarrow N$), medium ($-N$), or thick market ($\uparrow N$). I investigate two different simulation lengths: 5000 and 10000 respectively short ($\downarrow T$) or long ($\uparrow T$). Given the arrival rate, agents are expected to arrive 5 times in $\downarrow T$ and 10 times $\uparrow T$. Finally, I consider two different levels of adverse selection: 2×10^{-2} mean reversion with 5×10^6 fundamental shock variance and 2×10^{-4} mean reversion with 5×10^4 fundamental shock variance, respectively ($\downarrow A$) or ($\uparrow A$). The two adverse selection settings have roughly equal a priori final fundamental variance ($\downarrow A$ has 1.26×10^8 variance versus 1.23×10^8 in $\uparrow A$) in the long time horizon ($\uparrow T$) setting. These two settings differ in how the adverse selection changes over the time horizon. With high mean reversion there is less adverse selection at the beginning of the simulation, when most fundamental variation is damped by mean reversion. An environment is the combination of these three parameter settings, so I describe an environment in shorthand as the combination of the associated symbols. For instance, $\uparrow N \downarrow T \downarrow A$ is the environment with 238 agents and time horizon 5000, with a fundamental that has mean reversion of 0.02 and shock variance of 5×10^6.

The ZI strategy space defined in Section 2.2 is continuous, not necessarily compact, and not differentiable with respect to utility. I discretize the strategy space in order to make these games amenable to normal-form analysis. I want a set of discrete strategies such that equilibria in the discrete game have low regret in the continuous game. How to guarantee this property is unknown, but I nevertheless select what I consider a reasonable set of discrete strategies by starting with a set of feasible strategies, and then adding best responses to initial equilibria from a significantly larger set, as described in Section 3.1. The strategies I ultimately consider for equilibrium analysis are listed in Table 5.1.

Table 5.1: Strategies considered for equilibrium analysis. See Section 2.2 for parameter details. The first strategy is *truth-telling*, and the last strategy is effectively *fill-or-kill* where agents demand between 29 and 30 surplus.

$\underline{\boldsymbol{d}}$	0	0	0	0	0	0	0	0	125	250	500	500	1000	29000
$\overline{\boldsymbol{d}}$	0	125	250	500	1000	1500	2500	2500	1000	500	1000	1000	2000	30000
$\boldsymbol{\eta}$	1	1	1	1	1	0.6	0.4	1	0.4	1	0.4	1	0.4	0.001

In order to test the effect of tailoring strategies to the call mechanism, I also found equilibria in environments where agents have access to a select number of call strategies in addition to the regular background strategies. In the same way I chose a discrete set of standard strategies, I chose this set of call market strategies by looking for best responses from existing equilibria. The additional call market specific strategies are listed in Table 5.2. lists the three strategies agents had additional access to when finding equilibria with call strategies. All of these strategies exhibit decreasing levels of shading as the clear approaches, that is, as agents trade in a call market with better information, they shade less.

Table 5.2: Additional strategies considered for equilibrium analysis with call market strategies. See Section 5.3.1 for parameter details.

$\underline{\boldsymbol{d}}_{\boldsymbol{i}}$	250	29000	39000
$\overline{\boldsymbol{d}}_{\boldsymbol{i}}$	1250	30000	40000
$\underline{\boldsymbol{d}}_{\boldsymbol{f}}$	0	20000	30000
$\overline{\boldsymbol{d}}_{\boldsymbol{f}}$	1000	21000	31000
$\boldsymbol{\eta}$	0.4	0.001	0.001

5.4 Analysis

In order to understand how clearing interval affects market efficiency, I first demonstrate the effects when agents are non-strategic—specifically, when agents are truth-telling. I measure efficiency by fraction of optimal social welfare: average total trader surplus divided by the expected maximum possible trader surplus for the corresponding environment. This appropriately gives more weight to sampled instances with higher achievable surplus.

When the clearing interval is sufficiently small, I expect efficiency—as well as most other statistics—to be almost identical for call and CDA markets. The market types still will not be identical due to their different pricing schemes, and the discreteness of time. When the clearing interval is sufficiently large, I expect efficiency to degrade, because traders lack sufficient opportunities to trade. I also expect that at some point between almost-continuous clearing and a single clear, the efficiency of the call market will exceed CDA efficiency. The optimal clearing interval would ideally balance improvement in matching from aggregation over time, and provision to traders of ample opportunities to trade. In the words of Haas and Zoican (2016), when the clearing interval is short, the friction due to trader impatience is minimized, when the clearing interval is long, the friction due to imperfect information is minimized, and ideally at some intermediate interval, the two are balanced.

I simulated environment $-N{\uparrow}T{\downarrow}A$ ten thousand times for each clearing interval when agents submit truthful bids. Figure 5.1 shows the results of that simulation, which matches my expectations. The call market is generally more efficient than the CDA, as the clearing interval increases, the agents see the allocation benefit of batched clearing, until the point that infrequent trading opportunities dominate the benefit of batching, and there is a sharp decline in efficiency. Even at a clearing interval of 1 there is still a slight effect of batching as multiple agents may arrive at the same time interval. If the call market cleared after every interaction the incentives would still be different because the call market prices at the midpoint of trades not at the incumbent price.

Agents are unlikely to behave truthfully in real markets, and the model I investigate is too complex to support analytical solutions. I compute approximate empirical equilibria using the simulation-based process described in Chapter III—reducing the number of players down to four—on the environments discussed in Section 5.3.2. Figure 5.2 contains plots of clearing interval versus equilibrium efficiency for the six environments with low adverse selection. Since environments may have multiple equilibria, I plot a range of efficiencies for all of the equilibria found at a given clearing interval. In each CDA environment, I found exactly one equilibrium. These methods are not perfect, and occasionally produce *artifacts*—for instance, the dip in Figure 5.2d—which may be attributable to discretizations imposed in modeling. Despite the few artifacts, the general trends are unmistakable. The trend in efficiency closely matches what was expected and what we see in Figure 5.1, but to a lesser extent. Call markets provide less efficiency improvement over CDAs when agents are in equilibrium than when agents are truth telling. This decrease in efficiency improvement

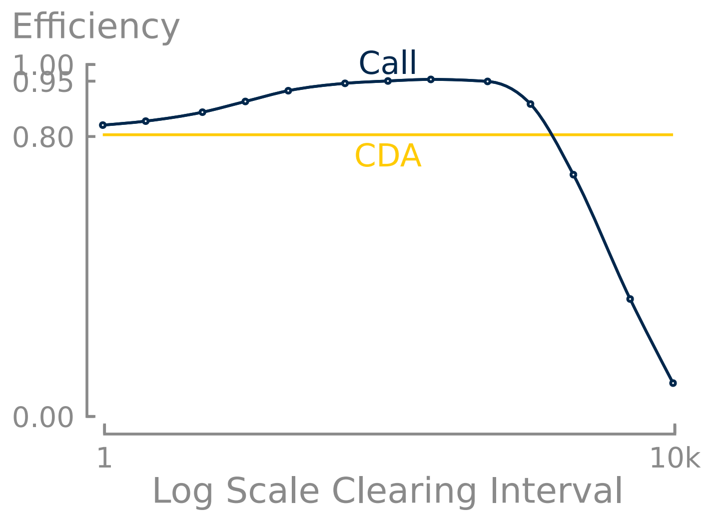

Figure 5.1: Call market efficiency in the $-N{\uparrow}T{\downarrow}A$ environment as a function of clearing interval when agents are truthful. In this environment, there is a large region of clearing intervals where call market efficiency is near optimal, and significantly higher than CDA market efficiency. Call market efficiency is slightly higher even when clearing every time step because multiple agents can arrive in a single time step and still benefit from batching.

is likely due in part to the surprising equilibrium efficiency of CDAs (Chapter IV).

As clearing interval increases, call markets go from being quite similar to CDAs to being significantly more efficient. In most of the environments the optimal call market mechanism reduces inefficiency by half when compared to the CDA. This trend is stronger as the number of agents increases, which I attribute partially to effective batching that does not happen when the markets are thinner. This suggests that frequent call markets would actually benefit thick markets more than thin, matching the result of Du and Zhu (2017). It also suggests that given an ideal market setting for call markets, the specific choice of clearing interval frequency is less important to efficiency improvement that using a call market at all.

There are other measures of market quality besides allocative efficiency, and the nature of this model allows me to compute arbitrary market statistics in equilibrium. Price discovery is a general measure of how closely the market represents the true price of a security; I calculate it as the *root-mean-square deviation* (RMSD) of the transaction prices relative to the unbiased estimate of the final fundamental at the time of transaction. A market with *worse* price discovery would have a *larger* RMSD. Call markets present a fundamental tension between common and private information aggregation: as the call interval gets longer, orders in the market represent more outdated fundamental information and may be less accurate, but as the call interval

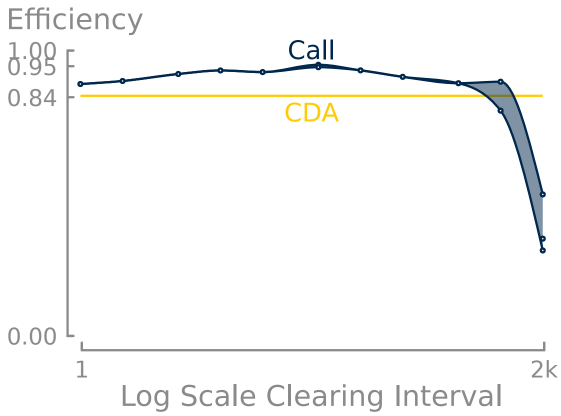

(a) Environment $\uparrow N \uparrow T \downarrow A$

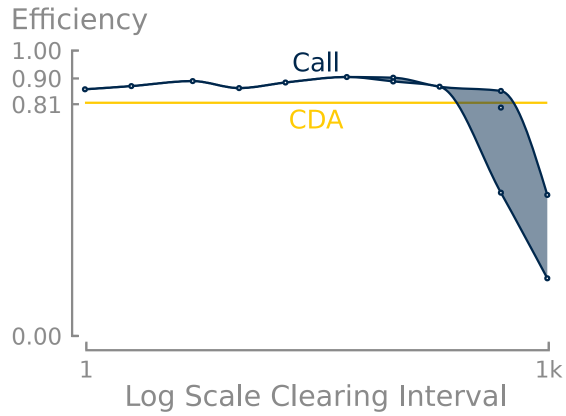

(b) Environment $\uparrow N \downarrow T \downarrow A$

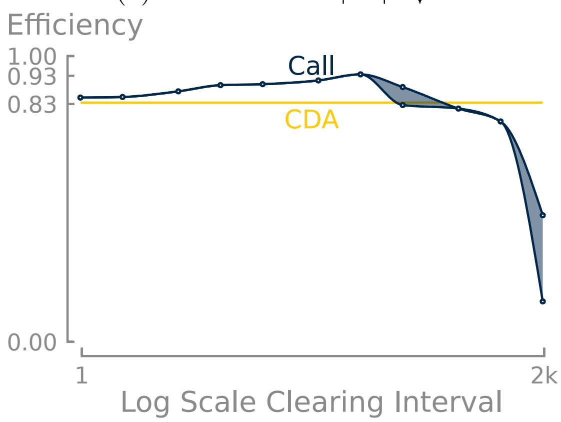

(c) Environment $-N \uparrow T \downarrow A$

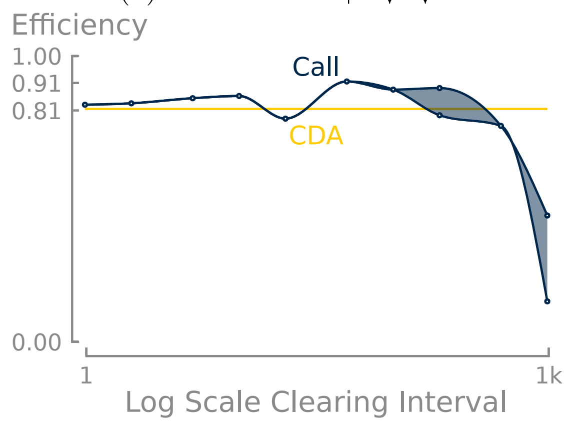

(d) Environment $-N \downarrow T \downarrow A$

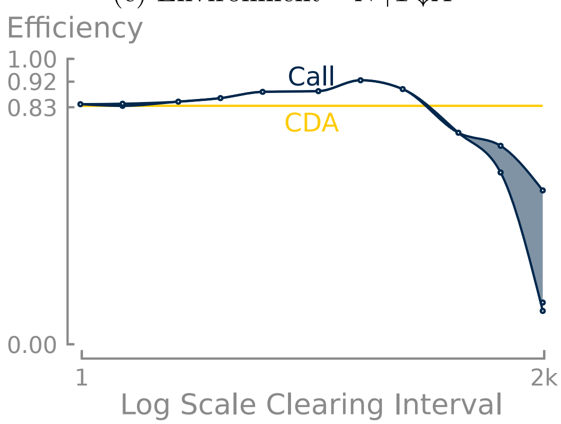

(e) Environment $\downarrow N \uparrow T \downarrow A$

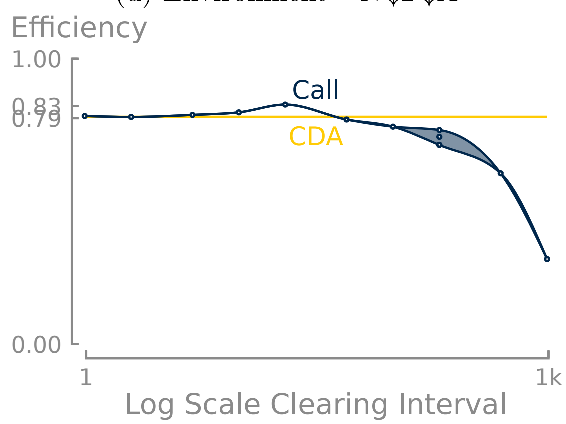

(f) Environment $\downarrow N \downarrow T \downarrow A$

Figure 5.2: Equilibrium efficiency as a function of clearing interval in six environments with low adverse selection. Rows from the top represent environments with 238, 58, and 25 agents respectively. The left column has a time horizon of 10000, while the right has a time horizon of 5000. Games where I found multiple equilibria have their efficiency represented as a range. These plots show the expected tradeoff between the benefits of batching and harm of limited execution opportunities. As the markets become thicker, the range of clearing intervals that are near optimal significantly increases.

gets longer more orders are present in each clear, and so more private information is aggregated in each transaction. Figure 5.3 presents the median RMSD for every low adverse-selection environment. CDAs exhibit better price discovery in every environment except for thick markets, and as for efficiency, when the clearing interval is long, price discovery significantly degrades. The price discovery improvement in thick markets likely arises because batching large numbers of orders helps to reduce the noise caused by private valuations, while still ensuring that orders represent recent information.

A common measure of liquidity is the market *spread*. The spread is the difference between the highest outstanding buy order and the lowest outstanding sell order. The size of the spread roughly corresponds to the transaction cost for impatient traders, as they must meet the quote to trade immediately. Since my model does not contain a dealer or market maker, there are times when spread is not defined, and I therefore report the median spread over all time. Increasing clearing interval should reduce the spread as agents have more protection from mispriced orders, since transactions happen at the midpoint of all orders instead of the incumbent price, and are therefore willing to price closer to their beliefs. This trend is present in Figure 5.4, which shows the equilibrium median spread for every low adverse-selection environment. The spread decreases as the clearing interval increases, but the benefit disappears as the clearing interval becomes prohibitively long.

My results indicate that long clearing intervals can result in low efficiency, poor price discovery, and large spreads, and the transition to this regime of poor market performance can be very fast. My results also indicate that appropriate clearing interval selection in thick markets with low adverse selection can significantly improve performance over CDAs. However, these results hold for strategies intended for use in CDAs. How robust are these results when agents can condition their strategy on their location in the clearing interval? Figure 5.5 shows the difference between equilibrium market statistics when call market agents have access to call market strategies. 82% of equilibria and 90% of games had equilibria where at least one call market strategy was in support, which indicates that it can be effective to develop customized strategies for call markets. However, despite incorporating some call market strategies, qualitative market performance in equilibrium is essentially the same.

My analysis to this point has considered only environments with low adverse-selection. Figures 5.6–5.8 show the same set of experiments for varying numbers of agents with a long time horizon and large adverse selection. I find that in high adverse-selection environments, the optimally efficient clearing interval approaches a

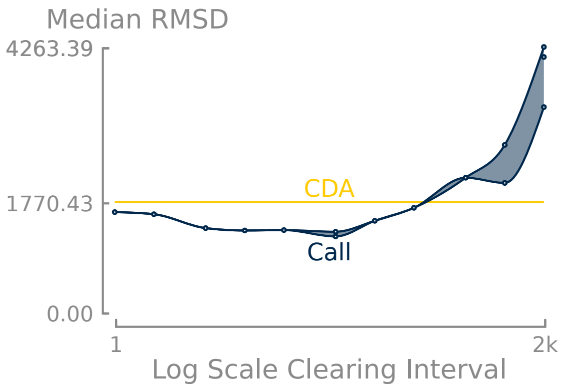

(a) Environment $\uparrow N \uparrow T \downarrow A$

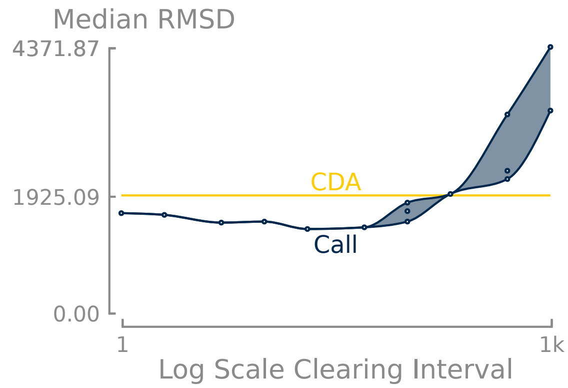

(b) Environment $\uparrow N \downarrow T \downarrow A$

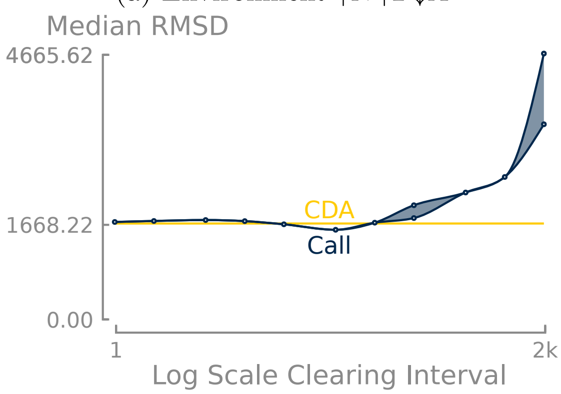

(c) Environment $-N \uparrow T \downarrow A$

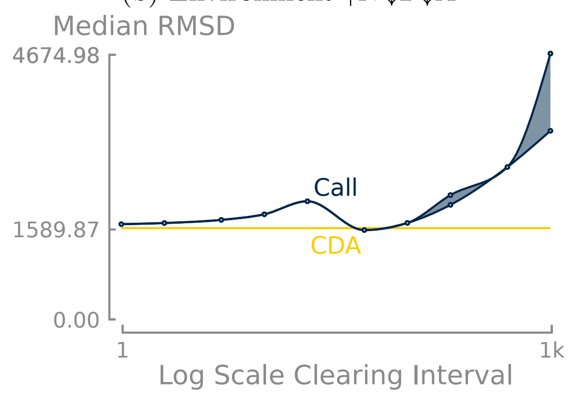

(d) Environment $-N \downarrow T \downarrow A$

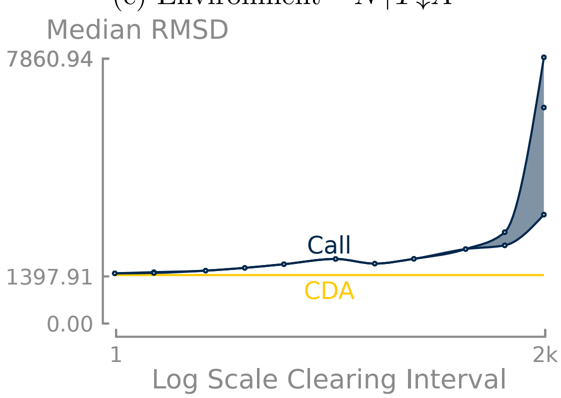

(e) Environment $\downarrow N \uparrow T \downarrow A$

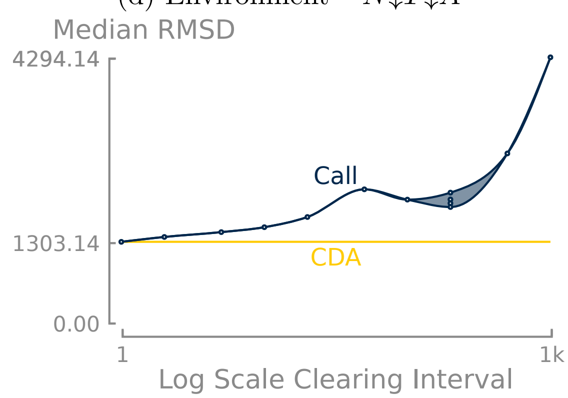

(f) Environment $\downarrow N \downarrow T \downarrow A$

Figure 5.3: Equilibrium price discovery as a function of clearing interval in low adverse-selection environments. Rows from the top represent environments with 238, 58, and 25 agents respectively. The left column has a time horizon of 10000, while the right has a time horizon of 5000. Higher RMSD indicates worse price discovery. Call markets improve price discovery only in thick markets with relatively short clearing intervals.

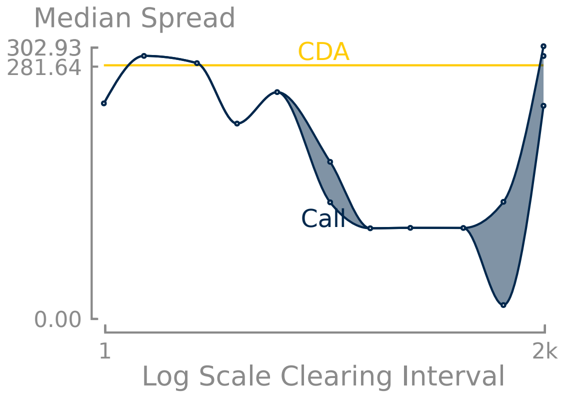

(a) Environment $\uparrow N \uparrow T \downarrow A$

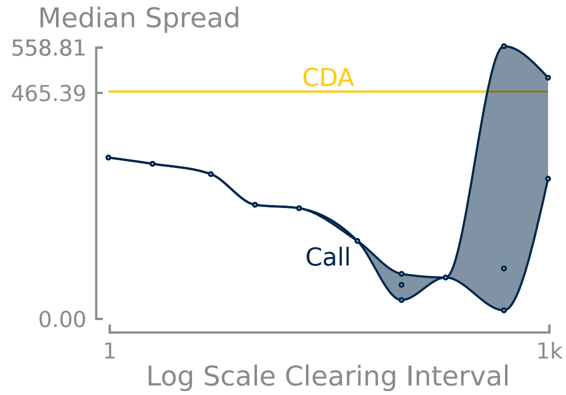

(b) Environment $\uparrow N \downarrow T \downarrow A$

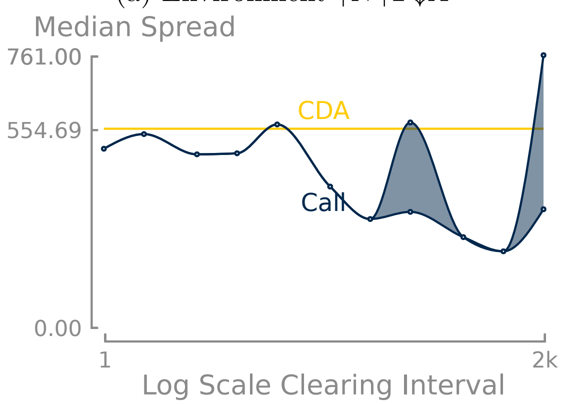

(c) Environment $-N \uparrow T \downarrow A$

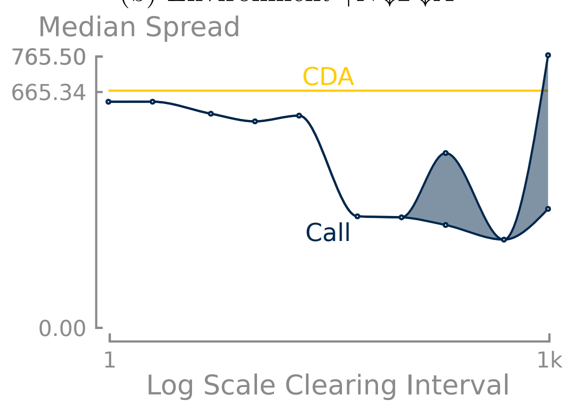

(d) Environment $-N \downarrow T \downarrow A$

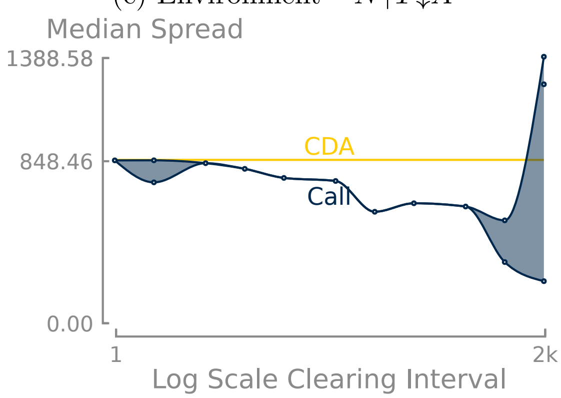

(e) Environment $\downarrow N \uparrow T \downarrow A$

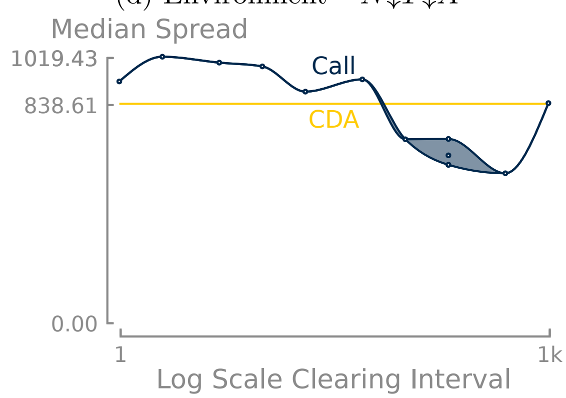

(f) Environment $\downarrow N \downarrow T \downarrow A$

Figure 5.4: Equilibrium median spreads as a function of clearing interval in low adverse-selection environments. Rows from the top are for environments with 238, 58, and 25 agents respectively; columns from the left represent time horizons of 10000 and 5000 respectively. Following intuition, longer clearing intervals decrease the spread; however, excessively long clearing intervals increase the spread.

CDA. Even when the number of agents is increased, short clearing intervals match, but do not exceed CDA efficiency. Thick markets do not restore price discovery as they do with low adverse selection, but the spread is generally tighter.

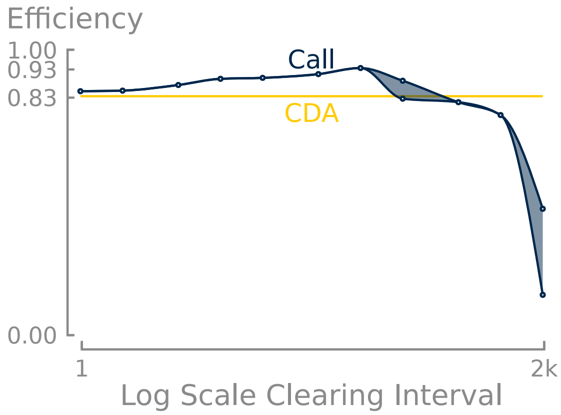

(a) Efficiency: Only CDA strategies

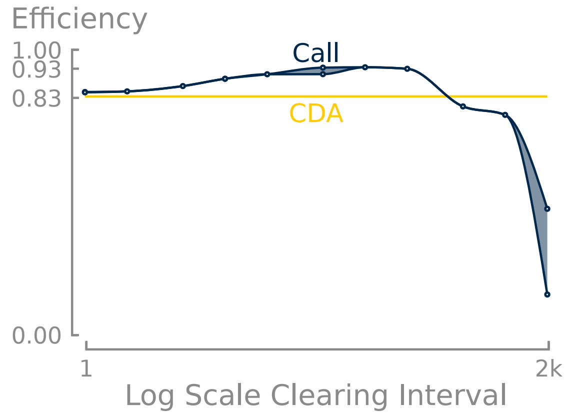

(b) Efficiency: CDA and call strategies

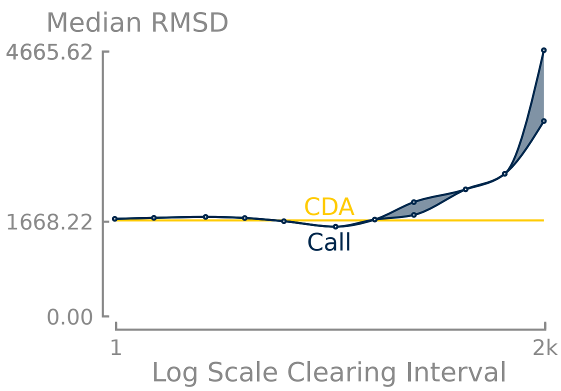

(c) Price Discovery: Only CDA strategies

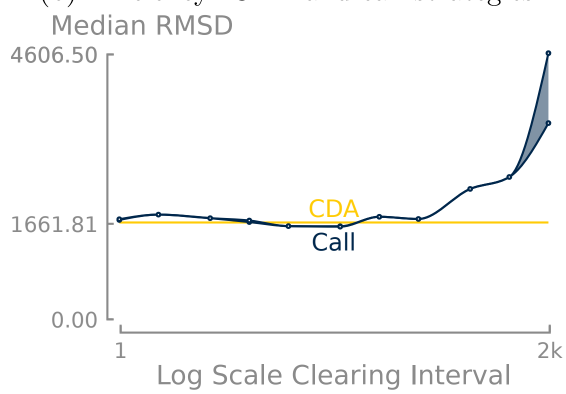

(d) Price Discovery: CDA and call strategies

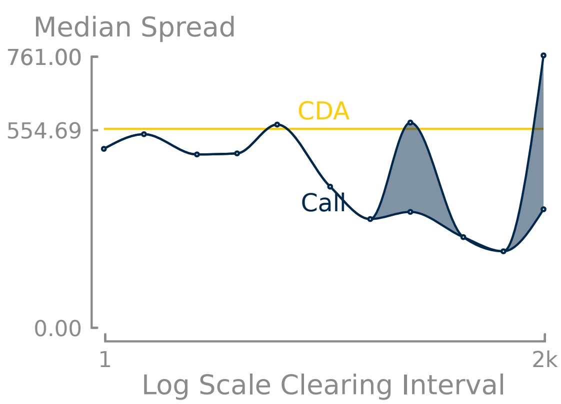

(e) Spread: Only CDA strategies

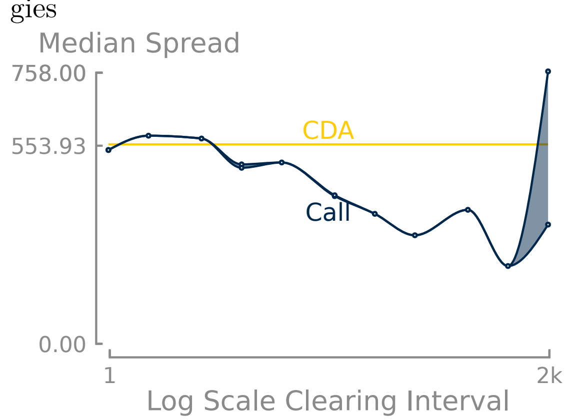

(f) Spread: CDA and call strategies

Figure 5.5: How equilibrium statistics change when agents have access to call market specific strategies in environment $-N{\uparrow}T{\downarrow}A$. The left column contains equilibria when agents have access only to CDA strategies; the right column contains equilibria when agents additionally have access to call market specific strategies. Call market strategies have weak effect on the qualitative trends in market quality.

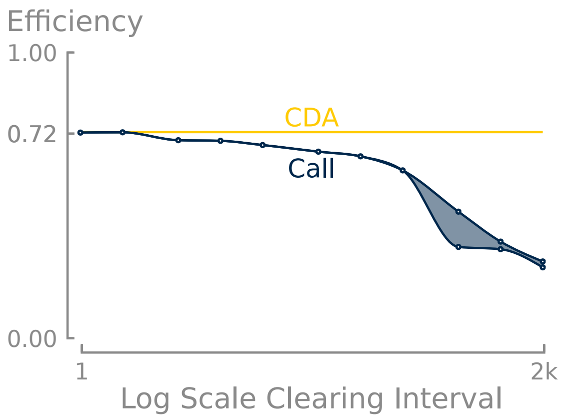

(a) Environment $-N\uparrow T\uparrow A$

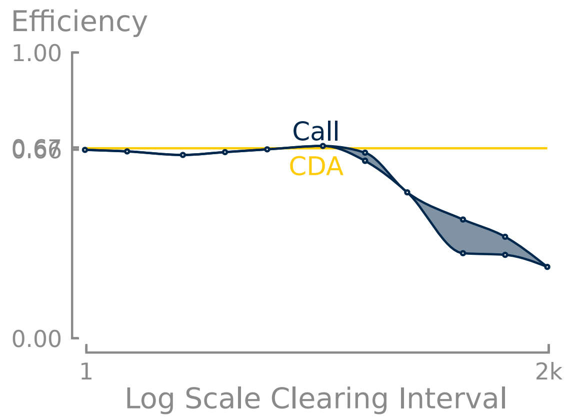

(b) Environment $\uparrow N\uparrow T\uparrow A$

Figure 5.6: Equilibrium efficiencies in two environments with high adverse selection. Even in a thick market (Figure b), short clearing intervals are on par with a CDA, but never surpass it.

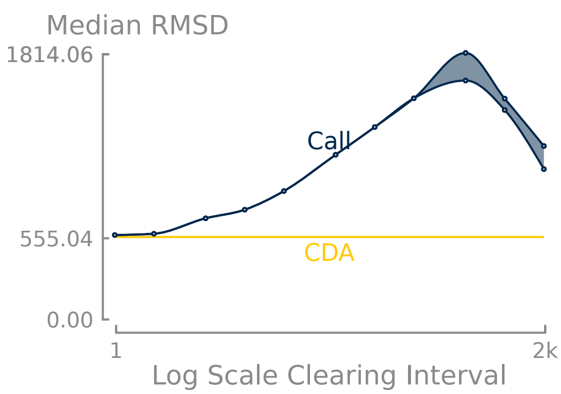

(a) Environment $-N\uparrow T\uparrow A$

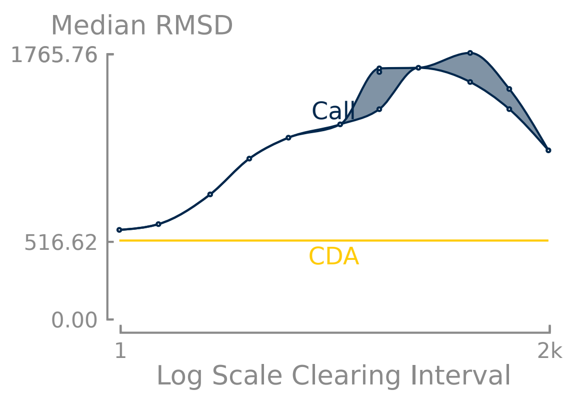

(b) Environment $\uparrow N\uparrow T\uparrow A$

Figure 5.7: Equilibrium price discovery with high adverse selection. Unlike with low adverse selection, price discovery in call markets with high adverse selection is consistently worse.

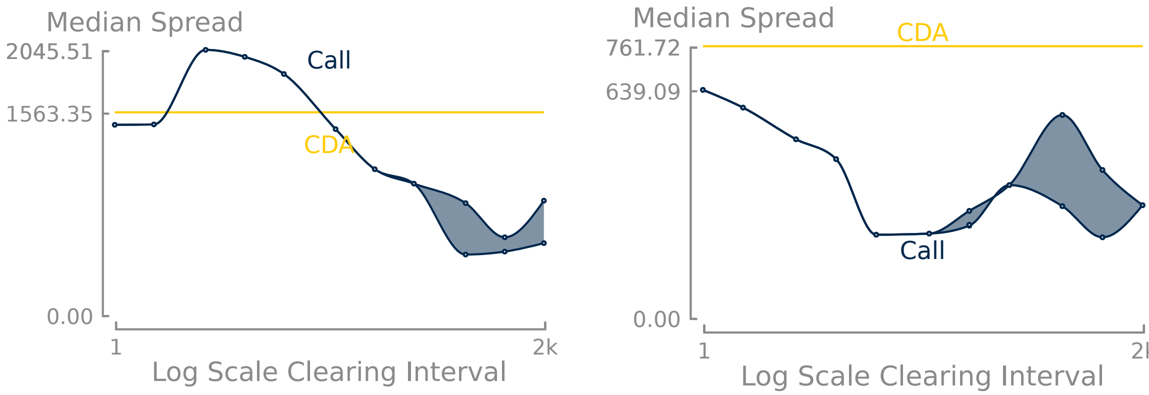

(a) Environment $-N\uparrow T\uparrow A$ (b) Environment $\uparrow N\uparrow T\uparrow A$

Figure 5.8: Equilibrium spread with high adverse selection. Spreads are reduced at large clearing intervals, even with high adverse selection.

5.5 Conclusion

I employed empirical mechanism design to investigate the effects of clearing interval on market statistics. Given the significant difference between equilibrium trading behavior in a call market and in a CDA, I believe analyses of strategic interaction, like the results presented here, are necessary to characterize expected behavior in call markets where little historic data exists. My analysis has produced a number of unique insights.

In a first exploration of agent strategies geared toward call markets, I find evidence suggesting that even simple consideration of the call interval can provide gains. However, while these strategies appear in equilibrium, they do not qualitatively change equilibrium market quality in my experiments. Are the lack of qualitative differences due to insufficient exploration of call market specific strategies, or does effective call market trading not need a significantly different strategy space? I leave this question for future work.

More to the point, there are a number of practical revelations about frequent call markets that I demonstrate. Environments with high adverse selection may be served more efficiently by CDAs than by frequent call markets. Otherwise, call markets tend to have some clearing interval region where they operate more efficiently, and this region is significantly wider in thick markets. Thick markets also universally saw improved spreads and price discovery than their thinner counterparts. These results roughly parallel those of Du and Zhu (2017), who also found long optimal clearing intervals in thick markets with low adverse selection. My results suggest that contrary to what might seem natural, the markets best suited to frequent call markets are thick with low adverse selection.

More generally, my methodology adapts readily to other metrics of interest, other market designs, and other agent behaviors, and provides a unique perspective on complex problems where models will always be stylized. For example, analyzing the effectiveness of random or volume-based clearing and other market-quality metrics can be easily factored into this model.

CHAPTER VI

Trend-Following Trading Strategies and Financial Market Stability

6.1 Introduction

What are the potential effects of trading automation on financial market stability? This is an obviously important question to answer in a time when trading activity is increasingly controlled by algorithms, interacting through electronic markets. Unprecedented information response times, autonomous operation, use of machine learning and other adaptive techniques, and ability to proliferate novel strategies at scale are all reasons to be concerned that algorithmic trading may produce dynamic behavior qualitatively different from what arises in trading under direct human control.

The question is also an extremely difficult one to answer. Though we already have significant experience with algorithmic trading, the available data can provide only limited support, for many reasons:

1. Algorithms operate at fine time scales over many markets (due to venue fragmentation or strategies operating across multiple securities), and so getting a consolidated view of behavior and cause-effect relations is challenging.

2. Trading strategies are proprietary and typically well-controlled secrets, limiting public knowledge about the specific methods in active use.

3. Market data reflects order actions and information flows, but does not directly reveal strategies.

4. New strategies are developed and deployed all the time.

5. Technology and market rules likewise are evolving.

6. The primary concern about stability relates to extreme but rare events, including possibility for severe reactions to conditions that have not yet been experienced and are thus not manifest in data.

6.1.1 Extreme Financial Market Events

The paradigmatic event of interest is a market crash, as exemplified by historical events of 1929 and 1987, and more recently the 2010 Flash Crash. The Flash Crash in particular has received a great deal of attention and analysis from regulators as well as academic researchers (Aldrich, Grundfest, and Laughlin 2016; Kirilenko et al. 2017; Menkveld and Yueshen 2017). As a result, there is a good overall understanding of the sequence of activity in that episode, and some insight into its causes and surrounding phenomena. Despite all this analysis, certain aspects remain quite mysterious, particularly regarding the unanswered question: could this happen again tomorrow? Many have suggested a positive answer—indeed observing that less intense flash crashes have happened, so actually the more salient questions would be under what circumstances would such crashes be likely to occur, and how could they be prevented or mitigated.

Stability concerns are particularly acute for high-frequency trading (HFT) strategies, due to the fine time scale at which they operate. Precisely because they adapt to changing conditions faster than any human could intervene, the potential for unintended reactions to rare events looms large.

Johnson et al. (2013) claim to find evidence of such HFT interactions in financial market data, which they term "ultrafast extreme events" (UEEs). Specifically, they document 18,520 instances in a five-year period of relatively large crashes or spikes in specific equity prices, lasting 1.5 seconds or less. The brief duration rules out human reactions in these price movements, and suggests the possibility of "new machine ecology" (or as Cartlidge and Cliff (2013) call it, a "robot phase transition") where algorithms interact in a series of responses, potentially producing novel price dynamics. However, Golub, Keane, and Poon (2012) shed doubt on the machine-ecology explanation of UEEs, demonstrating that these UEEs are actually better explained as the result of Intermarket Sweep Orders (ISOs), a kind of special order type introduced in the wake of market fragmentation. ISOs are exempt from the usual rules of order routing between exchanges, and can cause a mini crash or spike due to a large order on one single exchange.

The apparent confusion in interpreting these UEEs underscores the difficulty of characterizing algorithmic behavior in financial markets. It is highly likely that significant bot-only interactions do exist. Identifying their signatures and understanding their effects remain open problems for research. Effective monitoring of markets by regulators requires a fundamental detection capability, and ability to predict destabilizing conditions is prerequisite for designing market rules and interventions to promote stability.

6.1.2 Agent-Based Simulation

My simulations are *agent-based*, in that they directly model the decision behavior of trading agents interacting through market mechanisms, implemented at a level preserving key microstructural features. There is an extensive literature on agent-based modeling for generation of extreme events like flash crashes (Lee, Cheng, and Koh 2011; Paddrik et al. 2012). Bookstaber (2012) even argues that agent-based models are particularly suited to studies of financial stability, given the complexity of environments and interactions that raise stability issues.

A common criticism of agent-based modeling is that the results may be sensitive to details of the simulation implementation, and in particular the choices of agent behavior. In the present context, the fact that *some* configuration of agent behaviors could lead to extreme events in financial markets is not very conclusive, as these behaviors may themselves be highly implausible. I directly address this concern using the methodology presented in Chapter III.

To capture the causal pathways that may engender such instability, a financial market simulation must implement in detail the information structure and transmission by which agents' actions may influence subsequent actions by others. I specifically adopt a recent model (Wang and Wellman 2017) that limits the amount of information agents have about the fundamental price, thereby incentivizing them to incorporate market information in their strategies. This structure gives rise to a market that is more susceptible to shocks. I characterize stability by examining the impact of a market shock, one modeled as a large sell order, on the price of the security with and without trend followers. I find the presence of trend followers has weak effects on markets under normal conditions, but completely changes the nature of their instability. Without trend followers agents rely heavily on market information, which leads to market crashes with a long recovery time. When trend followers are present they disincentivize the use of market information, but amplify the impact of a market crash. The net result of introducing trend follows is a more significant market crash that recovers very quickly.

I also demonstrate a novel methodology to efficiently compute equilibria in games with an arbitrary probability of market shock. Equilibria with trend followers and endogenous shocks are more resistant to shocks that happen more frequently, but the presence of market shocks does not prevent trend following, keeping them unstable.

6.2 Related Work

Much of the research analyzing stability comes from empirical finance and econometric literature. Cetorelli et al. (2007) define a stable market as "one that can endure shocks to supply or demand without collapsing—that is, without experiencing surging (or wildly oscillating) prices". They focus their attention on centralization's effect on stability, and model their shock as a large supplier leaving the market. Their conclusions are mixed, finding both positive and negative relationships between centralization and stability. Baur and Schulze (2009) classify market shocks as systematic (frequent but small) and systemic (infrequent but large). Instead of directly characterizing the results of a shock, they focus on the change in impact between normal and extreme market conditions. Stable markets respond identically in both, whereas unstable markets exhibit a worse response during extreme conditions.

Empirical studies like these provide useful insights, but are not capable of addressing counterfactual questions that get at the causal factors underlying market stability. Agent-based simulation does support counterfactual comparison, and indeed has been extensively applied in this area. Lux and Marchesi (1999) describe a dynamic market model where agents alternate between being fundamentalists and effective trend followers. This simple model is expressive enough to give rise to large fluctuations in price despite low fluctuations in information and independent volatility, indicating a key relation between trend following and market instability. Hommes and Wagener (2009) use models with prescribed agent types and analyze the stability of the system as the parameters of strategies change. Their analysis finds bifurcation points where small levels of trend following are tolerable, but significant trend following leads to instability. A dynamic limit-order model of trend following—similar to the one presented in this paper—was introduced by Chiarella and Iori (2002) and extended by LeBaron and Yamamoto (2008). The latter paper demonstrates a connection between statistical observations from financial markets and the degree of trader imitation in their model. Lee, Cheng, and Koh (2010) propose a model where various aspects of market quality decrease as the number of trend following agents increase. Feng et al. (2012) use both an agent based and a purely stochastic model to recreate statistical phenomena found in transaction data. Technical trader imitation, i.e. similar behavior by non-fundamental traders, was necessary to exhibit fat tails in return distributions among other properties. These papers focused on statistical stability properties of markets with technical traders.

More recently, several authors have used agent based models to predict counterfactual responses to market shocks. Bookstaber and Paddrik (2015) present a model with agents that demand liquidity, supply liquidity, and market-make. They introduce a market shock by increasing the sell-side liquidity demand and analyze how that affects the overall liquidity of their model. The model proposed by Leal and Napoletano (2017) is closest to our own in that they precisely analyze the effect of a shock in a limit-order model. They study the effects of proposed policies on the volatility, frequency of crashes, and recovery time post-crash. Adding frictions to cancellations, such as minimum resting times or cancellation fees, proved to reduce volatility and frequency of crashes, but also lengthened recovery times.

My main point of departure from these prior agent-based studies is my imposition of a game-theoretic selection process on agent strategies.

6.3 Financial Market Simulator

I employ a version of the simulator presented in Chapter II, modified to capture imperfect fundamental observation, trend following, and market shocks. My simulations are composed of 66 background traders and optionally 486 trend followers. Table 6.1 lists the default configuration parameters, some of which are described later in this section.

Table 6.1: Default Environment Parameters

Parameter	**Variable**	**Value**
Background Arrival Rate	λ_B	2×10^{-3}
Fundamental Mean	$\bar{f}$	10^9
Fundamental Mean Reversion	r	0
Fundamental Shock Variance	σ_s^2	10^3
Max Position	$q_{\max}$	10
Observation Variance	σ_o^2	10^6
Private Value Variance	σ_{pv}^2	10^6
Trading Horizon	T	10^4
Trend Arrival Rate	λ_T	5×10^{-3}

6.3.1 Market Shocks

I model a shock to the market as a large incoming sell order, from a liquidating *shock agent* otherwise treated as outside the system. As a large sell order triggered the Flash Crash, this a natural form of shock for evaluating market stability.

In order to avoid having the shock agent directly set the price of trade during the shock, in our implementation the agent liquidates its holdings evenly throughout a given interval (Bertsimas and Lo 1998). Specifically, the shock agent arrives at every time step and tries to to reduce its fraction of holdings to the proportion of liquidation time remaining by taking any, possibly all, outstanding buy orders. This is slightly different from the liquidation strategy used by traders blamed for instigating the Flash Crash, which chose the amount to liquidate proportional to the previous net trading volume.

6.3.2 Background Agents

Background agents arrive at the market and act with probability 2×10^{-3} every time step, giving them twenty opportunities to trade in expectation. This rate is typically sufficient to obtain much of the potential gains from trade, as competitive equilibrium with this valuation structure more than 95% of agents hold an absolute position less than five.[1] I adopt the information structure from Wang and Wellman (2017) as an attempt to account for the (realistic) prospect that agent actions can reveal meaningful information about common value components. At every arrival, agents observe a noisy signal drawn around the fundamental, formally $\tilde{f}_t = f_t + n_t, n_t \sim \mathcal{N}(0, \sigma_o^2)$. The observation variance (σ_o^2) is set to to 10^6 by default, roughly equal in magnitude to the expected variance of the fundamental between arrivals, 5×10^5.

A background agent assesses its expected valuation at the time of market entry t, using an unbiased estimate $\hat{f}_t$ of the current fundamental f_t. In all environments in this chapter, the fundamental mean reversion (r) is set to 0, making $\hat{f}_t$ an unbiased estimator of the final fundamental f_T. The estimate is updated every market entry using posterior inference from the new noisy fundamental observation ($\tilde{f}_t$) and the last estimate from time t'. As the Gaussian distribution is its own conjugate prior, and closed under linear transformations, an agent's belief about the fundamental at

[1] I estimate this by sampling populations from the valuation distribution, and solving for competitive equilibrium in the population instance. Of course background traders do not accrue all available gains due to inefficiencies of CDAs, including strategic effects—though shading per se can improve efficiency in CDAs (Chapter IV).

a specific time can also be represented compactly as a mean ($\hat{f}_t$) and variance ($\hat{\sigma}_t^2$).

$$\hat{f}_t = \frac{\sigma_o^2 \hat{f}_{t'} + \tilde{\sigma}_t^2 \tilde{f}_t}{\sigma_o^2 + \tilde{\sigma}_t^2}, \quad \hat{\sigma}_t^2 = \frac{\sigma_o^2 \tilde{\sigma}_t^2}{\sigma_o^2 + \tilde{\sigma}_t^2}, \quad \tilde{\sigma}_t^2 = (t - t')\sigma_s^2 + \hat{\sigma}_{t'}^2, \quad \hat{f}_0 = f_0, \quad \hat{\sigma}_0^2 = 0. \quad (6.1)$$

When the market forbids market orders, the threshold parameter (η) can be used as a means to submit so-called "fill-or-kill" orders, which are designed to trade immediately or not at all by leaving unfilled orders at an unreasonable price (described in Section 2.2). With the prospect of large market shocks, however, it is difficult to set a fixed price that is guaranteed to be unreasonable, and so I explicitly support fill-or-kill orders that are submitted only when they trade immediately.

6.3.3 Shock Transmission

In order for an agent to be significantly influenced by another's actions, and thus amplify a shock, it must be through some impact of those actions on the state of the world or the agent's beliefs. Consider two agents. Agent 1 may affect the world state by issuing a buy order, which will either get placed in the order book or transact with an existing sell order. If agent 2's values and belief are unchanged, then this state change would make it only more likely that agent 2's action in response will tend to counteract agent 1's. It is only if agent 1 changes agent 2's beliefs or values in the direction of its own action that we might expect to see propagation or amplification.

Background agent valuations have private and common components. An agent's action may reveal something about its own private component, but in most contexts this is weak information relative to the common prior. Of greater potential significance is how the action reflects the agent's information about the common fundamental. This potential hinges on imperfect observation of the common value component. If each agent observes the *true* fundamental immediately prior to submitting orders ($\sigma_o^2 = 0$), as was the case in the previous chapters, the Markovian nature of fundamental evolution (2.1) renders irrelevant any prior observations of the fundamental, and thus agents gain no information about common value by considering others' actions.

Figure 6.1a shows an example time series when a shock is introduced to a market where agents do not consider consider the actions of others. When I introduce a price shock (formally specified in Section 6.3.1), there is an impact to the order book, but agent strategies are oblivious to transaction prices and so continue submitting orders based on observed fundamental signals.

When agents receive *noisy* information about the fundamental ($\sigma_o^2 > 0$) other-agent behavior leaks relevant information about the fundamental value. I extend the

(a) Agents do not regard one another's actions (i.e., ignore price history). In this case, the large shock impacts the order book, but agents values still stay close to the true fundamental, and prices return to normal almost immediately after the shock.

(b) Agents account for transaction prices in their estimate of the fundamental value, rendering the market more vulnerable to shocks. In this case, the shock influences agents' fundamental estimates, which deepens the effect and makes it more persistent.

Figure 6.1: Effect of market shocks in the absence and presence of price-conditional strategies. The y-axis ("Deviation") represents deviation from the initial fundamental value. In both plots, the yellow line represents the fundamental value of the security, the gray line is the transaction price series without a market shock, and the blue line is the same price series, but with a market shock from time 5,000 to time 5,500. Due to common randomness, the price series are identical prior to the market shock.

background trading strategy to allow agents to use past market data to enhance their estimate of the fundamental. Specifically, the agents use price information from past transactions as if they were drawn from a Gaussian distribution centered around the fundamental at the time of transaction. Whenever an agent observes a transaction it runs the update step in (6.1) but it uses the transaction price instead of $\tilde{f}_t$ and an assumed strategic variance $\sigma_p^2 \in [0, \infty]$ instead of the fundamental observation variance σ_o^2. Setting σ_p^2 to infinity is equivalent to ignoring transaction prices and using the normal background strategy. When σ_p^2 is finite, agents still use fundamental information in their value estimate, but weighted with previous transaction prices. When agents choose to use market information in estimating their value, it exposes the market to shocks, as agents will use asset pricing during a shock to update their belief about the fundamental value. This effect is present in Figure 6.1b, where everything else is as in Figure 6.1a, but background agents are using market information.

6.3.4 Trend Following

I next consider the possibility of exploiting price trends, which has the potential to exacerbate price shocks. As noted above, when the fundamental shifts significantly, subsequently arriving background agents are more likely to transact with orders on the side opposite the direction of change. Due to the limited arrivals of background agents, their stale orders can persist for some time. As a result, market price lags behind the fundamental during large shifts inducing a correlation between long price trends and fundamental mispricing. This correlation can be seen in the yellow line in Figure 6.2, where I plot the difference between the current transaction price and current fundamental value with respect to the length of a monotonic price trend. Formally, the plot represents the following:

$$\text{Fundamental Difference}(r) = \mathbb{E}[f_i - p_i \mid p_i \geq \cdots \geq p_{i-\ell+1}], \tag{6.2}$$

where ℓ is the length of a trend, p_i is the i^{th} transaction price, and f_i is the fundamental value at the time of the i^{th} transaction.

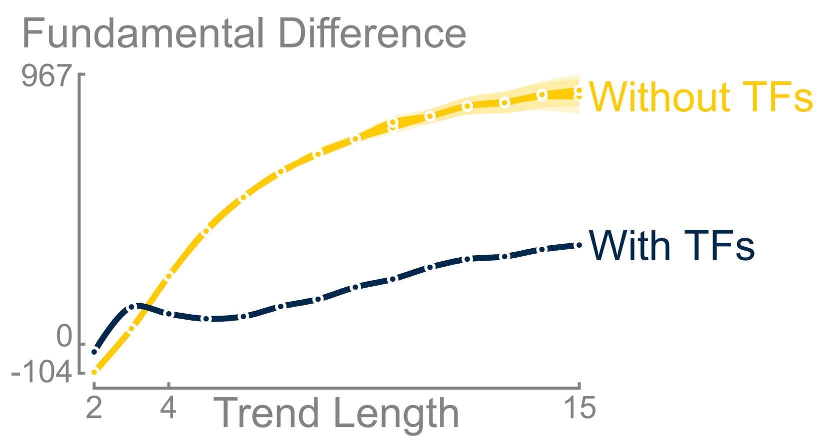

Figure 6.2: Price trends correlate with mispricing of the fundamental value. A trend of length ℓ is present if each of the last $\ell-1$ transactions were at greater prices than the preceding transaction. "Fundamental Difference" is the average difference between the current fundamental value and the current transaction price (Equation (6.2)), here plotted in yellow as a function of trend length in equilibrium without the presence of a market shock. Introducing trend followers (blue), acting on trends at least four long ($L = 4$) in equilibrium, helps rectify this inefficiency and brings the expected fundamental at transaction time closer to transaction prices.

Agent-based models of trend-following agents have been studied in prior work (Leal and Napoletano 2017; Lee, Cheng, and Koh 2010), but none are close fits to our market environment due to the way they match orders. Chiarella and Iori (2002) present some analysis of a model with trend followers, that has similar structure to ours—with a limit-order book and single unit orders. Their trend followers linearly interpolate a fundamental value from past transactions prices. I introduce a simple trend-following agent, designed specifically to exploit the mispricing during fundamental shifts, highlighted in Figure 6.2, so that I can measure its effect on market stability and evaluate whether this behavior is sustainable in equilibrium. Trend followers arrive in the market with the same geometric process as background agents, but at a higher rate, 2.5 times that of background traders or 5×10^{-3}. Upon arrival, a trend-following agent checks whether the past L transaction prices are monotone. If they are monotonically increasing, the agent first takes the outstanding ask at price p, which exposes the next lowest outstanding ask of q (interpreted as infinite if the quote is missing). The trend follower then resubmits a new sell order at $\min\{p + \Pi, \max\{q - 1, p\}\}$, attempting to make up to Π profit, but guaranteeing they place the best ask. In this way, the trend follower actively continues (follows) price trends. If the new order does not transact within time t_{exp} it is withdrawn and the agent accepts the final liquidation price for the unit it bought. Trend followers perform symmetrically when the trend is decreasing. They have no private value, and do not get any information about the fundamental value; they act solely on the basis of the correlation between price trends and fundamental shifts.

The blue line in Figure 6.2 illustrates that introducing trend followers (with $L = 4$), removes some of the mispricing, bringing the trend-conditional transaction prices closer to the fundamental on average. When I exogenously add trend followers to an environment with a shock, as seen in Figure 6.3, their effect is stark. The trend followers slightly help the market respond to the change in fundamental before the shock hits, and afterward exacerbate the security's precipitous decline.

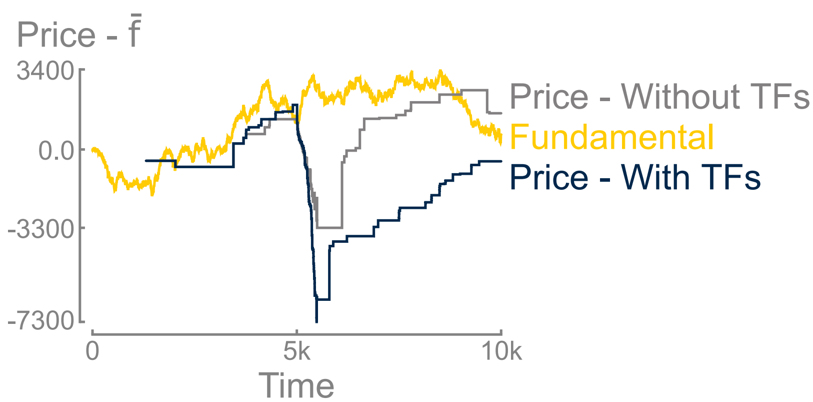

Figure 6.3: Trend following exacerbates the market's response to a price shock. The yellow line represents the fundamental value of the security. The gray line represents transaction prices during a market shock without trend followers, and the blue line represents the same transaction prices, but with trend followers.

6.4 Market Shocks on Equilibrium Behavior

For considerations of stability, a key question is how widely agents will adopt strategies that make significant use of transaction information. Will they still adopt those strategies in the presence of trend followers, and if so, will there still be a significant impact due to market shock? I can address this question with the EGTA methodology discussed in Chapter III. I systematically simulate strategy profiles from the market game defined in Section 6.3 where agents can play only strategies from the restricted strategy space defined in Tables 6.2 and 6.4, in order to identify symmetric mixed-strategy Nash equilibria. I sampled each selected profile 30,000 times to consider it *evaluated*, and used DPR to reduce background agents down to 6 and trend followers down to 2. I stopped requiring that restricted games be explored if they had more than three strategies in support for games without trend followers, or four strategies in support across both roles for games with trend followers.

Table 6.2: Background agent strategies considered in systematic equilibrium search. $\underline{d}$ and $\overline{d}$ are bounds on the amount of extra surplus a background agent demands with each trade. η is what fraction of the random surplus an agent is willing to accept if it gets immediate execution; when η is fill-or-kill (FOK), agents submit an order only if it will transact for the desired surplus. σ_p^2 is the assumed transaction price variance around the fundamental; when it is infinite, the strategy ignores transaction price information. In the default environment, the empirical variance of price around the fundamental is 3×10^5.

$\underline{d}$	$\overline{d}$	η	σ_p^2
500	1000	0.8	10^6
500	1000	0.8	10^9
500	1000	0.8	∞
1000	1500	0.8	10^6
1000	1500	0.8	10^9
1000	1500	0.8	∞
0	500	1	∞
30	30	FOK	10^6
30	30	FOK	10^9
30	30	FOK	∞
90	90	FOK	10^6

My first experiment entails a quantification of the response to an exogenous market shock when agents are in Nash equilibrium. I compute equilibria with and without trend followers and for each equilibrium I verified there is low regret with respect to the larger strategy sets of Tables 6.3 and 6.5. Every confirmed equilibrium includes

Table 6.3: Larger set of background agent strategies considered for potential deviation from found equilibria

$\underline{d}$	$\overline{d}$	η
500	1000	0.8
1000	1500	0.8
0	500	1
30	30	FOK
90	90	FOK

$\times$

σ_p^2
∞
10^3
10^6
10^9

Table 6.4: Trend follower strategies considered in systematic equilibrium search. L is the length of transaction price trend necessary to act. A setting $L = 2$ would be trivially satisfied by any two transactions. Π is the profit sought, conditioned on seeing a trend. t_{exp} indicates how long a trend follower waits after placing a trend order before withdrawing it.

L	Π	t_{exp}
4	100	50
5	100	20

Table 6.5: Larger set of trend follower strategies considered for potential deviation from found equilibria

L
4
5
8
10

$\times$

Π
20
50
100

$\times$

t_{exp}
20
50
100

positive support for a strategy that uses price information, but when trend followers are present, the only background strategies that use past price information are fill-or-kill. When I allow trend-following agents to exist, they are present in equilibrium, and background agents continued to use pricing information, although to a lesser degree.

I analyze the effect of trend followers with three characteristics of market performance: transaction price volatility, mispricing, and maximum fundamental difference (Max Difference). Transaction price volatility is the sample standard deviation of transaction prices. Common empirical measures of market stability tend to relate to volatility. Mispricing is the root-mean-squared difference (RMSD) between each transaction price, and the fundamental value at the time of transaction, first introduced in Section 5.4. Mispricing roughly corresponds to the area between the fundamental and transaction price time series. In the presence of a shock, mispricing also indicates the impact of the shock, accounting for both the significance of the price change and the duration. Maximum fundamental difference, is the largest absolute difference between a transaction price and the fundamental value at the time of transaction. This metric emphasizes the instantaneous magnitude of the shock ignoring the recovery time.

For each equilibrium I found in the environment with and without trend followers, I sampled ten thousand deviation profiles for each strategy to estimate the empirical regret in the restricted game, and the expected value of each metric over the equilibrium distribution. The empirical regrets, bootstrap confidence bounds, and regrets to the expanded strategy space are presented in the first rows of Table 6.6, where shock variance is 1000 and observation variance is 10^6.

Table 6.6: Empirical and bootstrap regrets for equilibria with and without trend followers in five environment settings. "Shock Variance" and "Observation Variance" are environment variables. "Trend Followers" indicates the presence (✓) or absence (✗) of trend followers. "Num" labels the equilibria. "Mean Regret" is the mean empirical regret obtained from taking i.i.d. samples from the equilibrium distribution. "95% Regret" is the bootstrapped 95% regret confidence interval calculated from the same samples as "Mean Regret." "Full Regret" is the DPR Deviation regret for deviating to the full strategy sets in Tables 6.3 and 6.5.

Shock Variance (σ_s^2)	Observation Variance (σ_o^2)	Trend Followers	Num	Mean Regret	95% Regret	Full Regret
1000	10^6	✗	1	129.20	206.46	0.00
			2	100.42	162.36	0.00
		✓	1	82.62	155.02	1.37
1000	10^5	✗	1	76.24	121.16	0.00
		✓	1	73.86	85.63	0.00
			2	77.58	96.54	4.08
1000	10^4	✗	1	117.67	166.95	7.89
			2	77.06	127.54	0.39
		✓	1	77.99	135.23	0.00
100	10^6	✗	1	9.01	26.12	1.96
			2	0.00	15.37	10.85
		✓	1	0.37	11.89	0.19
10	10^6	✗	1	0.00	8.94	3.15
		✓	1	0.08	0.16	6.70
			2	0.00	11.08	6.70

First, I investigate how the presence of trend followers affects volatility and price discovery in equilibrium. Figure 6.4 shows the volatility with and without trend followers. Trend followers appear to lower volatility, but the magnitude is small and could easily be an artifact of strategy set. Figure 6.5 shows the price discovery, measured as RMSD, in the market with and without trend followers. There is a small reduction in price discovery (larger RMSD) in the presence of trend followers, however similar to volatility the magnitude is not large enough to make any claims with confidence. The only market statistic that did change noticeably was the volume, an expected consequence of significantly more traders.

Figure 6.4: Trend followers slightly reduce volatility in equilibrium. Volatility is measured as the standard deviation of transaction prices, and in light of the forty times increase in trading volume with trend followers, a small reduction in volatility is not unexpected. There are two bars without trend followers because I found two equilibria in that environment. These equilibria have similar makeup and, therefore, have similar volatility.

Figure 6.5: Trend followers slightly increase mispricing in equilibrium. Mispricing is measured as the root mean squared difference (RMSD) between the fundamental and transaction prices. There are two bars without trend followers because two equilibria were found in this environment, however these equilibria are very similar which is why their volatilities are similar.

I now investigate how markets with and without trend followers respond in the presence of an exogenous market shock. I expose each market to various levels of market shocks, up to an agent liquidating 100 units (the average volume in the market without trend followers). The liquidation rate of the shock is held constant between magnitudes, so a shock of 10 units takes twice as long as a shock of 5 units. Figure 6.6 shows the maximum difference between a transaction price and the current fundamental price in markets without and without trend followers for various levels of market shock. For shocks under 20 orders, the markets behave roughly identical. The shock takes deep orders in the order book, but otherwise has a similar level of effect between markets without and without trend followers. Once the shock goes beyond 20 units, the trend followers amplify the effect. Large market shocks are roughly twice as severe in markets with trend followers.

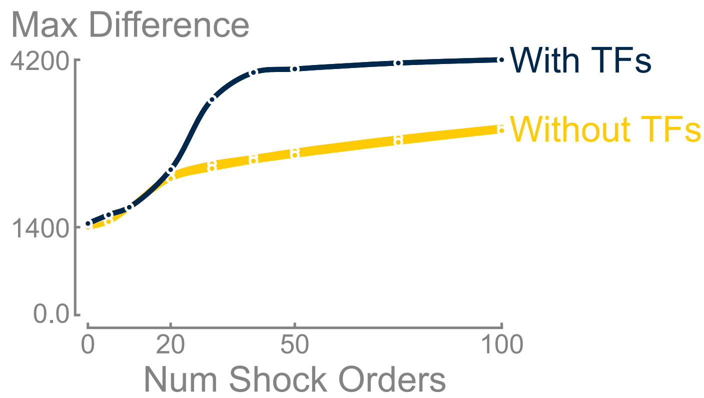

Figure 6.6: Trend followers amplify the decline in price in the presence of large market shocks. When market shocks are small markets with and without trend followers behave almost identically, but once a shock destroys the order book, trend followers amplify it. “Max Difference” is the largest absolute difference between a transaction price and the fundamental price at the time of transaction.

Figure 6.7 shows the mispricing in markets with and without trend followers for various levels of market shock. Just like maximum difference, for small shocks, markets with and without trend followers behave almost identically, but when the shocks get large enough, something strange happens, the mispricing is actually worse in the market without trend followers. The cause of this strange phenomenon lies in the strategies that are played in equilibrium. Without trend following, price making agents use transaction information in both equilibria. This makes non-trend-following environments prone to shock like Figure 6.1b. Since only fill-or-kill strategies use price information with trend followers, market shocks do not get amplified by background agents. The result is that shocks in trend-follower equilibrium look more like Figure 6.8, a market with trend followers but where background agents don’t use price information. Shocks are more severe in terms of maximum fundamental difference, but are also much quicker to recover.

Trend followers change the fundamental style of the crash that results from a large shock order. Without trend followers, the crash is slow and changes fundamental traders beliefs. There is a period where fundamental traders slowly observe new information and adjust their beliefs back towards the true fundamental. When trend followers enter the market, the incentives shift to reduce the use of transaction information. Crashes are now much more severe, but recover just as quickly, taking a form more qualitatively similar to the Flash Crash.

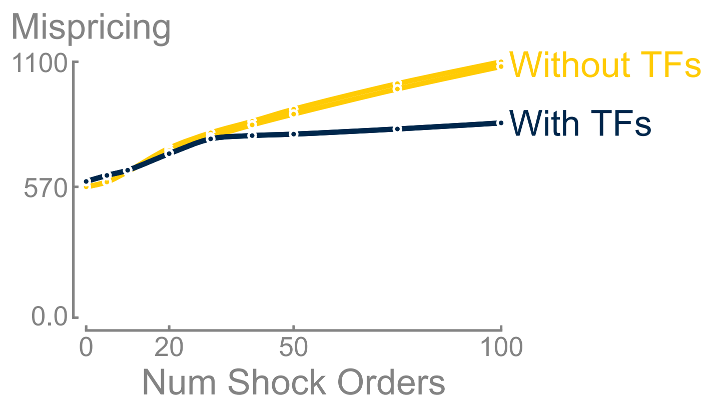

Figure 6.7: Trend followers improve the mispricing (RMSD) in response to market shocks. They crowd out background traders, removing their incentive to learn from historic prices. The recovery post shock is faster because background agents beliefs do not need to reset to the true fundamental. Trend followers also follow trends in either direction, helping increase the magnitude of the shock, but also helping its recovery. These two factors contribute to a lower shock mispricing in the presence of trend followers. The mispricing with 0 shock orders are the same values as presented in Figure 6.5.

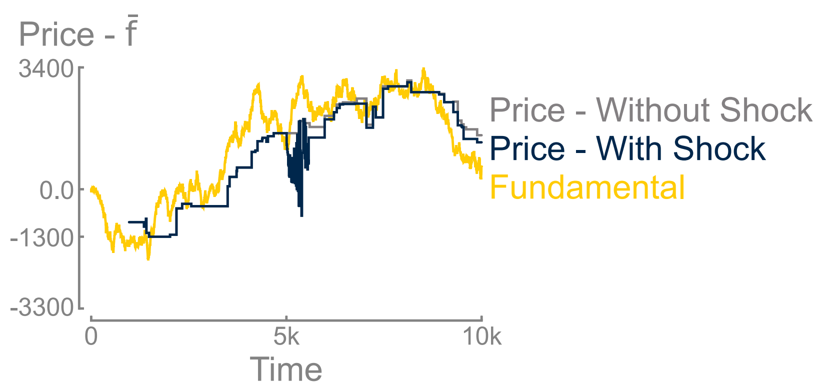

Figure 6.8: Trend followers in markets where agents do not use transaction information produce instantaneous crashes, where price drops significantly, but recovers just as quickly. This style of market shock is what we actually see when equilibrating markets with trend followers, as background agents primarily ignore past price information.

6.5 Effects in Alternate Environments

The previous results characterize effect of trend followers on stability in one environment, but the viability of trend following is heavily predicated on the delayed response of background agents to fundamental shifts. I investigate the effect of market shocks in equilibrium as the fundamental volatility or the observation variance is reduced. Regrets for the equilibria found in these environments is listed in Table 6.6.

In order to prevent confusion between a market shock, instigated by a large sell order, and the fundamental shock variance σ_s^2, which controls the magnitude of fundamental shifts, I use the term "fundamental volatility" in the place of "fundamental shock variance." Lower fundamental volatility reduces the likelihood of large fundamental shifts, the source of trend follower profit, while lower observation variance lowers the incentive to use transaction price information and reduces the mispricing after a fundamental shift.

Figure 6.9 shows the maximum fundamental difference and RMSD for a 50 unit shock as the fundamental volatility is decreased. When fundamental volatility is reduced to 10, trend following is no longer profitable, so trend followers exit the market making it very stable in the presence of shock. The middle fundamental volatility produces shock response between the first two. The same qualitative effects are still present: trend followers significantly increase the instantaneous drop of the market shock and lower the RMSD. Each effect is damped as background agents can use more conservative strategies and trend followers have fewer opportunities for profit.

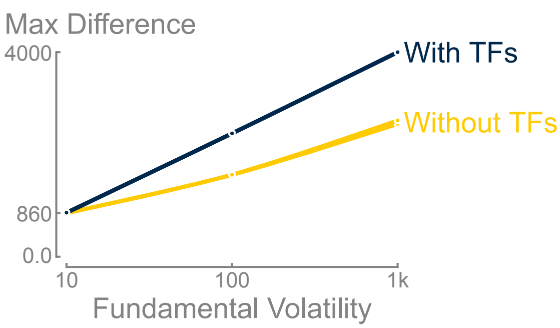

(a) Maximum fundamental difference in response to a 50 unit exogenous shock in equilibrium as a function of fundamental shock variance.

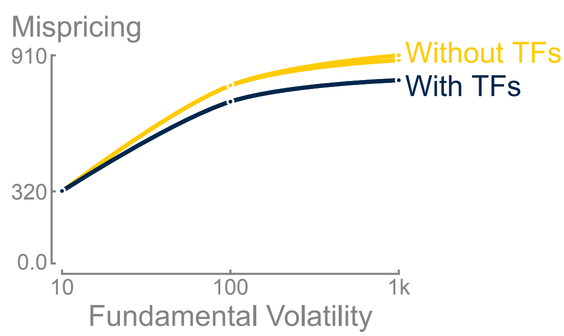

(b) Mispricing (RMSD) in response to a 50 unit exogenous shock in equilibrium as a function of fundamental shock variance.

Figure 6.9: Reducing fundamental shock variance, one part of the friction that trend followers exploit, increases market stability. When fundamental shock variance is 10, trend following is no longer profitable, and trend followers leave the market making it very stable in the presence of shocks. Between both extremes the qualitative effects are the same, but exist to a lesser magnitude. These results serve to validate that the model behavior matches expectations, and that the magnitude of shock impact is dependent on the specific environment, but that the qualitative effect is preserved.

Figure 6.10 shows the maximum fundamental difference and mispricing for a 50 unit shock as the fundamental volatility is decreased. When observation variance is reduced, trend following is no longer profitable, so trend followers exit the market.

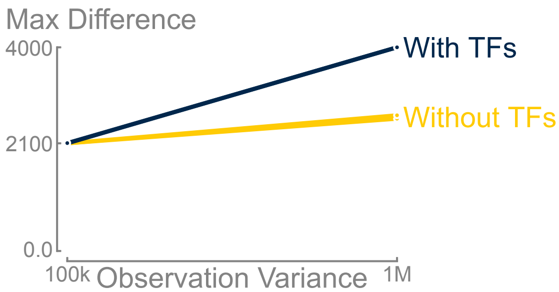

(a) Maximum fundamental difference in response to a 50 unit exogenous shock in equilibrium as a function of observation variance.

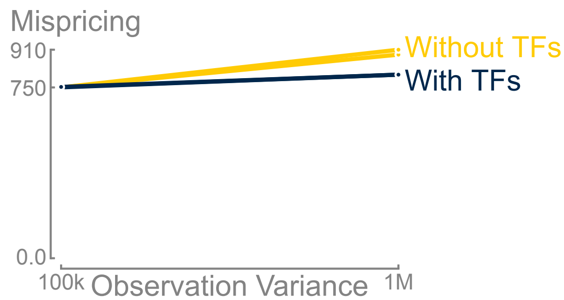

(b) Mispricing (RMSD) in response to a 50 unit exogenous shock in equilibrium as a function of observation variance.

Figure 6.10: Reducing observation variance removes incentives to trend follow, increasing market stability.

These results are expected given the results from Section 6.4, and so serve to validate the consistency of the model and that the qualitative effects are less sensitive to market parameter settings—as long as there is enough fundamental variation and observation variance for trend followers to exist.

6.6 Anticipating Market Shocks

The results from Sections 6.4 and 6.5 are derived assuming an unanticipated market shock. These sections discuss market effects when agents assume market shocks never happen, but how do these results change when agents anticipate a low probability shock?

In principle, I want the equilibrium strategies when the environment has some probability of shock. These results could be naively obtained by simply defining a new environment where a shock happens with some probability, and reapplying the methodology from Chapter III for every anticipated probability of shock I wished to measure. However, each of these games has an important structure, the payoffs are a convex combination of the payoffs in the environment without shocks and the environment with shocks. If agents anticipate a shock with probability t, then the expected payoffs in that environment are $t(\text{payoff with shock}) + (1-t)(\text{payoff without shock})$. Instead of defining new environments for every probability of anticipated shock, an environment can evaluate payoffs by simply weighting the evaluated payoffs from each of these games. This would still require running the profile search methodology for each probability t I wished to evaluate.

I refer to the general class of games, whose payoffs are the convex combination of two other games, as *mixture games*. I present a novel technique to compute nearby empirical mixture game equilibria requiring the minimal number of profile evaluations. I use this technique in combination with the profile search algorithm form Chapter III to compute equilibria for a wide range of anticipated shock probabilities.

6.6.1 Nearby Mixture Game Equilibria

To compute nearby mixture game equilibria, I construct an ODE that dictates how an equilibrium changes with respect to the mixture between two games. Let $\mathbf{m}_t$ be an equilibrium of the t mixture game, defined as as the game where payoffs come from game A a $1-t$ fraction of the time, and game B otherwise. In general, there may be more than one equilibrium, so $\mathbf{m}_t$ is set valued function where the set in nonempty. However, we can defined a choice function that selects a specific equilibrium at each t. There are often specific choice functions that are continuous and differentiable for most values of t as demonstrated in the examples in Section 6.6.2. We assume that $\mathbf{m}_t$ is the result of one such choice function so that it is continuous and differentiable at t and compute its derivative under this assumption. Note that for a specific equilibrium, there may not exist a choice function such that $\mathbf{m}_t$ is continuous and differentiable at t. Since equilibria are easy to verify, if the ODE ever returns an invalid equilibrium, we know this assumption was violated.

We also assume that $\mathbf{m}_t$ has *full support*, without loss of generality. Note that if $\mathbf{m}_t$ is an equilibrium in the unrestricted game, we can restrict the strategy set to only strategies that are in support. $\mathbf{m}_t$ is an equilibrium in the restricted game and we can verify if it is an equilibrium of the unrestricted game by checking the payoffs to unilaterally deviating; if deviating outside the support is ever beneficial then $\mathbf{m}_t$ is not an equilibrium. Define $a_s(\mathbf{m}_t)$ and $b_s(\mathbf{m}_t)$ to be the payoff of unilaterally deviating to strategy s when all other agents play according to $\mathbf{m}_t$ in games A and B respectively, where $s \in \mathcal{S}, |\mathcal{S}| = S$. Both $a_s(\mathbf{m}_t)$ and $b_s(\mathbf{m}_t)$ are the expectations of multinomial distributions and are therefore continuous and differentiable w.r.t. $\mathbf{m}_t$. Dot notation is used to indicate the derivative with respect to t, and ∇ is used to represent Jacobian.

From the definition of the mixture game, the payoff for unilaterally deviating to strategy s when all other agents play according to $\mathbf{m}_t$ in the t mixture game, $p_{ts}(\mathbf{m}_t)$, is:

$$p_{ts}(\mathbf{m}_t) = (1 - t)a_s(\mathbf{m}_t) + tb_s(\mathbf{m}_t).$$

Since a_s and b_s are continuous and differentiable, and we assume $\mathbf{m}_t$ is continuous and differentiable locally, p_s is also locally continuous and differentiable. Taking the derivative with respect to t yields:

$$\begin{aligned}\dot{p}_{ts}(\mathbf{m}_t) &= (1-t)\nabla^\top_{\mathbf{m}_t} a_s(\mathbf{m}_t)\dot{\mathbf{m}}_t + t\nabla^\top_{\mathbf{m}_t} b_s(\mathbf{m}_t)\dot{\mathbf{m}}_t - a_s(\mathbf{m}_t) + b_s(\mathbf{m}_t), \\ &= \mathbf{f}_s^\top(\mathbf{m}_t, t)\dot{\mathbf{m}}_t - g_s(\mathbf{m}_t),\end{aligned}$$

where

$$g_s(\mathbf{m}_t) = a_s(\mathbf{m}_t) - b_s(\mathbf{m}_t), \qquad \mathbf{f}_s(\mathbf{m}_t, t) = (1-t)\nabla_{\mathbf{m}_t} a_s(\mathbf{m}_t) + t\nabla_{\mathbf{m}_t} b_s(\mathbf{m}_t).$$

All strategies in support of a Nash equilibrium must have equal expected payoffs, therefore $p_{tr}(\mathbf{m}_t) = p_{ts}(\mathbf{m}_t)\, \forall t \in [0,1] \wedge r, s \in \mathcal{S}$, since $\mathbf{m}_t$ is an equilibrium of the t mixture game and all strategies are assumed to be in support. Since the payoff to unilaterally deviating is also continuous and differentiable at t, the derivatives are also equal, $\dot{p}_{tr}(\mathbf{m}_t) = \dot{p}_{ts}(\mathbf{m}_t)$. Substituting our derivation of the derivatives yields:

$$\left(\mathbf{f}_r^\top(\mathbf{m}_t, t) - \mathbf{f}_s^\top(\mathbf{m}_t, t)\right)\dot{\mathbf{m}}_t = g_r(\mathbf{m}_t) - g_s(\mathbf{m}_t).$$

From here, $S-1$ independent pairs of strategies can be chosen to generate $S-1$ independent equalities; one simple choice of pairs is $\{(s,i) \mid i \in \mathcal{S} \setminus \{s\}\}$. One final equality constrains the equilibrium mixture to remain in the simplex, $\sum_i \dot{m}_{ti} = 0$ or $\mathbf{1}^\top\dot{\mathbf{m}}_t = 0$. If I take the strategies as indices from 1 to S and use the pairs suggested, I can represent these S equations in matrix form,

$$\begin{aligned}\mathbf{F}(\mathbf{m}_t, t)\dot{\mathbf{m}}_t &= \mathbf{g}(\mathbf{m}_t) \\ \dot{\mathbf{m}}_t &= \mathbf{F}^{-1}(\mathbf{m}_t, t)\mathbf{g}(\mathbf{m}_t)\end{aligned}$$

where

$$\mathbf{F}(\mathbf{m}_t, t) = \begin{bmatrix} \mathbf{f}_1^\top(\mathbf{m}_t, t) - \mathbf{f}_2^\top(\mathbf{m}_t, t) \\ \vdots \\ \mathbf{f}_1^\top(\mathbf{m}_t, t) - \mathbf{f}_S^\top(\mathbf{m}_t, t) \\ \mathbf{1}^\top \end{bmatrix} \qquad \mathbf{g}(\mathbf{m}_t) = \begin{bmatrix} g_1(\mathbf{m}_t) - g_2(\mathbf{m}_t) \\ \vdots \\ g_1(\mathbf{m}_t) - g_S(\mathbf{m}_t) \\ 0 \end{bmatrix}$$

This final equation represents the derivative of the components of an equilibrium mixture with full support in a t mixture game. There are four possible events that can happen when solving this ODE to compute equilibria in a mixture game:

1. t becomes 0 or 1: The ODE is undefined as it is only valid for t in $[0, 1]$.

2. a strategy probability drops to zero: The zero probability strategy can be removed from the support and the ODE is still defined for the smaller support.

3. a strategy outside of support becomes a beneficial deviation: The ODE is no longer defined as $\mathbf{m}_t$ is not an equilibrium.

4. $\mathbf{F}$ becomes singular: the derivative may exist, but it is not the unique solution to the system of equations and it would therefore require other information to derive. For simplicity we just assume the ODE is no longer defined.

5. $\mathbf{m}_t$ is no longer an equilibrium: The ODE is no longer defined. This can happen when the equilibrium is not differentiable, but $\mathbf{F}$ is nonsingular.

Notably, computing the derivative only requires all payoff data in support of the equilibrium (to compute the deviation payoffs and their Jacobians) and the payoffs for unilaterally deviating outside the support to know if the equilibrium is still defined, for each game. This is also all of the data needed to verify an equilibrium, and is therefore minimal.

6.6.2 Mixture Game Equilibria Examples

In order to better understand the space of mixture game equilibria, and what their derivatives look like, I present mixture equilibria in two simple two-player two-strategy symmetric games. The game matrices are presented as row play payoffs.

Figure 6.11 shows a mixture game with a continuous set of equilibria when $t = \frac{1}{2}$. The mixture game transitions from a game where Down is the dominant strategy to a game where Up is the dominant strategy. When the mix is slightly biased towards either game, only the dominant strategy exists as an equilibrium. The derivative is a constant zero in either of these regions, and ceases to be valid once deviating to the other strategy becomes beneficial. When the games are evenly mixed, all distributions are equilibria, but there is no choice function such that an equilibrium is differentiable. Matrix $\mathbf{F}$ is singular at all of these points.

Figure 6.12 shows a discontinuity in the space of equilibria. The mixture game transitions from a game with one equilibrium of always play Up, to a game where either pure strategy and the mixture between the two are equilibria. When the games are evenly mixed, playing Down is a Nash equilibrium, but is not trembling hand stable. As the second game becomes dominant, multiple equilibria exist, each of which is locally continuous and differentiable w.r.t. t.

	U	D
U	0	0
D	1	1

(a) Row player payoffs when $t = 0$. Down is the dominant strategy.

	U	D
U	1	1
D	0	0

(b) Row player payoffs when $t = 1$. Up is the dominant strategy.

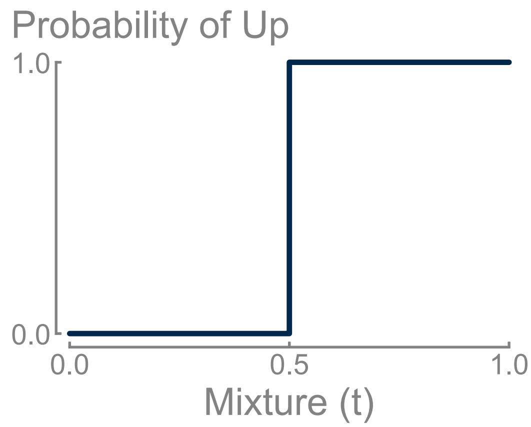

(c) Equilibria as a function of mixture fraction between games

Figure 6.11: Mixture equilibria where multiple equilibria exist. Matrix **F** is undefined when $t = 0.5$.

	U	D
U	1	2
D	0	1

(a) Row player payoffs when $t = 0$. Up is the only equilibrium.

	U	D
U	1	0
D	0	1

(b) Row player payoffs when $t = 1$. Up, Down, and picking uniformly randomly are all equilibria.

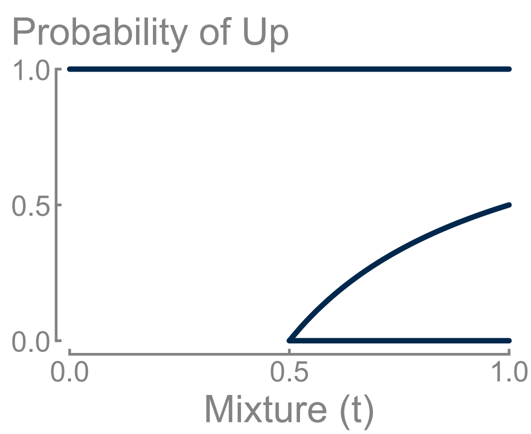

(c) Equilibria as a function of mixture fraction between games

Figure 6.12: Mixture equilibria where equilibria are discontinuous. As t approaches 0.5, two of the equilibria converge and then cease to exist.

6.6.3 Mixture Game Equilibria Tracing Technique

The ODE from Section 6.6.1 can be readily applied to compute an equilibrium for any value of the mixture parameter t. The technique I describe is recursive, and finds a continuous range of equilibria for every t within a range $[\underline{t}, \overline{t}]$. The technique uses numerical ODE solving to compute equilibria as a function of t. The numeric nature means there is no closed form solution for each equilibrium as a function of t, but instead a range of t for which the equilibrium is continuous is identified. Within that range, the ODE can be used to find an equilibrium value for any specific t.

1. Compute equilibria of the $\frac{1}{2}(\underline{t} + \overline{t}) = \tau$ mixture game using the technique from Chapter III. The mixture game can be computed by mixing profiles after sampling them to minimize profile evaluation.

2. Use the ODE from Section 6.6.1 on each τ game equilibrium to find an equilibrium for the large and smallest values of t possible. Let $\underline{\tau}$ ($\overline{\tau}$) represent the minimum (maximum) t with which an equilibrium was confirmed.

3. Recursively apply this procedure on the ranges $[\underline{t}, \underline{\tau}]$ and $[\overline{\tau}, \overline{t}]$, if they are not empty.

Starting with the range $[0, 1]$, this procedure will find a set of continuous equilibria as a function of t that cover $t \in [0, 1]$.

6.6.4 Results

I applied the technique from Section 6.6.3 to the environment from Section 6.4 to evaluate how the equilibria change as agents anticipate different probabilities of a 50 unit shock. I plot results with and without shock, even though agents are in equilibrium for the environment when shocks happen a certain fraction of the time. This allows comparison of market behavior between normal and extreme conditions while agents' prior over the probability of shock changes. Figure 6.13 shows the equilibrium maximum fundamental difference with and without a shock as the probability of endogenous shock increases. Trend following remains profitable even when shocks happen 98% of the time. Shocks provide a large influx of surplus as the shock agent trades below market value. Trend followers are net profitable despite losing money during the downward crash; they recoup enough of their losses on the rebound. In addition, as agents anticipate shocks more, the market becomes more resilient to their effects. As the probability of endogenous shock increases, background agents use transaction price information less, resulting in lower equilibrium maximum fundamental difference.

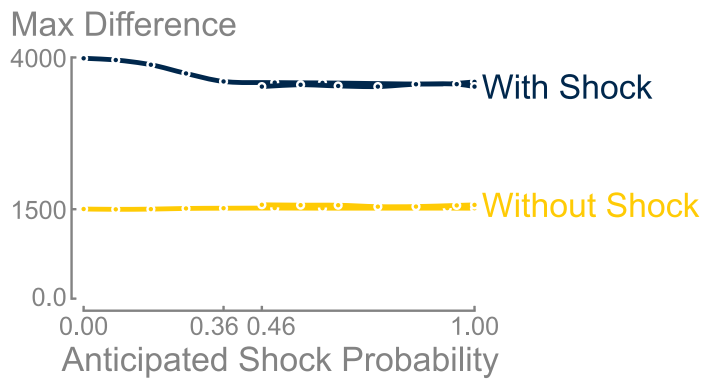

Figure 6.13: Maximum fundamental difference with and without market shocks as the equilibrium market shock probability changes. On the left side, market shocks are completely unanticipated, that is, agents are in a nash equilibrium when a shock never happens. On the right, agents anticipate a shock every simulation. A small amount of anticipation produces more stable markets in the presence of trend followers, but the result is small. Trend followers continue to exist, even when market shocks always happen. When the anticipated shock probability increases above 0.46 a new equilibrium exists, but its maximum difference is close to the other equilibrium.

Figure 6.14 presents the same style of plot, but showing RMSD instead of maximum fundamental difference. The RMSD is mostly constant as endogenous shock probability changes.

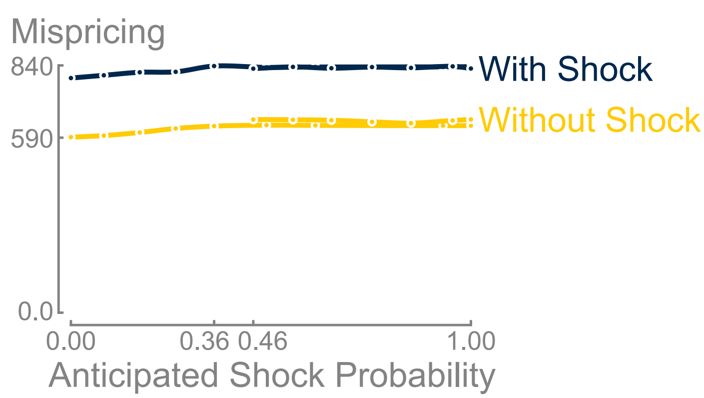

Figure 6.14: Mispricing (RMSD) as the probability of anticipated shock changes. Mispricing remains relatively constant even when agents highly anticipate a shock. Trend followers recoup most of their losses on the rebound, and so continue operating even when shocks are frequent.

6.7 Conclusion

The simple and realistic assumption that traders have imperfect information about a common (fundamental) value opens a path for these agents to seek information in each others' action, which in turn enables the propagation of market shocks—implemented as a large exogenous sell order. I studied the effects of such a large market shock when agents play according to approximate Nash equilibrium strategies. In this environment, market shocks have a low price impact with a slow recovery time. Agents' delayed market access allows for profitable equilibrium trend following, which changes market shocks to have high price impact but a fast recovery. Absent the presence of triggering shocks, these two response types have nearly identical volatility and price discovery. When agents anticipate a low probability of shock the stability of the market increases, but it does not remove the incentive for trend following or significantly alter the impact of the shock.

CHAPTER VII

Conclusion

Financial market activity is increasingly influenced by algorithms, accounting for over half of all trading volume in U.S. exchanges. Unprecedented information response times, autonomous operation, use of machine learning and other adaptive techniques, and ability to proliferate novel strategies at scale are all reasons to question whether algorithmic trading may produce dynamic behavior qualitatively different from what arises in trading under direct human control. Given the high level of competition between trading firms and the significant financial incentives to trading, it is reasonable consider the effect incentives have on the behavior of agents in financial markets.

In this dissertation I examined how market structure incentivizes agent behavior, and the effect of that behavior on market performance. I approached these analyses using agent-based models that capture the relevant microstructural aspects of financial markets and trading strategies of interest. Using agent-based models to study financial markets is not new, but my thesis builds on recent work applying empirical game-theoretic analysis to study agent based models in equilibrium. I explored three distinct case studies on the effect of market structure on performance. What follows is a summary of my contributions from the three case studies.

Shading and Efficiency in Limit-Order Markets (Chapter IV) In this Chapter I significantly extended a previous work showing notable equilibrium efficiency in continuous double auctions (CDAs) (Zhan and Friedman 2007). I extended the scope of the previous results in two significant ways. First, I improved upon the analysis techniques for finding Nash equilibria, including developing a system to more easily analyze games of this size (Section 3.2). More importantly, I also extended the market model from a very stylized model of a CDA, to a model with the significant informational, heterogeneous, and strategic complexities of modern financial markets. My research shows that even with the complexities of financial markets, CDAs are still remarkably efficient in equilibrium. This is true despite the fact that the ability to trade in both directions undoes the clear dissection of extramarginal and intramarginal inefficiency—the major justification for the effect. I also showed that in our models, market aspects that are thought to generally hurt efficiency, e.g. limited trader arrivals, improve the equilibrium efficiency gain relative to truth telling, and are therefore less harmful than may be naively thought.

Empirical Mechanism Design for Optimizing Clearing Interval in Frequent Call Markets (Chapter V) In this chapter, I evaluate the optimal setting for the clearing interval in a frequent call market using empirical mechanism design. This mechanism has been proposed as a way to prevent the deleterious effects of certain high-frequency trading strategies, but has also shown promise in increasing general trader welfare. Despite a large body of research suggesting why they should be implemented, very few studies have thoroughly discussed how the frequency should be set. My results indicate that while call markets provide helpful aggregation of private valuations, they hurt the aggregation of common information by giving equal importance to stale orders. In thick markets with low adverse selection, frequent call markets provide welfare improvement for a wide range of frequencies, but with high adverse selection, any call market may actually hurt welfare. I also provide the first implementation of call market specific strategies in an agent-based model. These strategies appear in equilibrium, but they do not change any qualitative aspects of my results. Their specification still serves as a baseline for further agent-based studies of call markets.

Trend-Following Trading Strategies and Financial Market Stability (Chapter VI) In my final study I examine the impact of introducing trend followers on financial market stability. I measure a market's stability as its response to a large exogenous sell order—a shock. When trend followers are present in the market, market shocks are more severe, producing a larger instantaneous mispricing. Markets with trend followers also recover from shocks faster than markets without trend followers. The faster recovery happens for two separate reasons: 1. trend followers help the recovery by also following the trend back to the fundamental price, and 2. trend followers remove the incentive for background agents to use past transaction price information, removing the delay for background agents to update their beliefs about the fundamental post shock. These results are robust to the amount of fundamental volatility, providing it is high enough for trend followers to be profitable.

This thesis expands upon the scalability of the underlying methodology, broadening the scope of problems EGTA can address. In it, I have advanced the understanding of three important structural questions facing modern financial markets by using agent-based modeling and game-theoretic strategy selection.

BIBLIOGRAPHY

BIBLIOGRAPHY

Akerlof, George A. (1970). "The Market for "Lemons": Quality Uncertainty and the Market Mechanism". In: *Quarterly Journal of Economics* 84.3, pp. 488–500.

Aldrich, Eric M., Joseph A. Grundfest, and Gregory Laughlin (2016). *The Flash Crash: A New Deconstruction*. Tech. rep. Available at SSRN, 2721922. University of California–Santa Cruz and Stanford University.

Baldauf, Markus and Joshua Mollner (2015). *High-Frequency Trade and Market Performance*. Tech. rep. Available at SSRN, 2674767. University of British Columbia and Northwestern University.

Baur, Dirk G. and Niels Schulze (2009). "Financial market stability: A test". In: *Journal of International Financial Markets, Institutions and Money* 19.3, pp. 506–519.

Bertsimas, Dimitris and Andrew W. Lo (1998). "Optimal control of execution costs". In: *Journal of Financial Markets* 1.1, pp. 1–50.

Biais, Bruno, Larry Glosten, and Chester Spatt (2005). "Market microstructure: A survey of microfoundations, empirical results, and policy implications". In: *Journal of Financial Markets* 8.2, pp. 217–264.

Bookstaber, Richard (2012). *Using Agent-Based Models for Analyzing Threats to Financial Stability*. Working Paper 3. Office of Financial Research.

Bookstaber, Richard and Mark Paddrik (2015). *An Agent-based Model for Crisis Liquidity Dynamics*. Tech. rep. US Treasury Office of Financial Research.

Budish, Eric, Peter Cramton, and John Shim (2014). "Implementation Details for Frequent Batch Auctions: Slowing Down Markets to the Blink of an Eye". In: *American Economic Review* 104.5, pp. 418–424.

— (2015). "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response". In: *Quarterly Journal of Economics* 130.4, pp. 1547–1621.

Cartlidge, John and Dave Cliff (2013). "Evidencing the 'robot phase transition' in human-agent experimental financial markets". In: *Fifth International Conference on Agents and Artificial Intelligence*. Barcelona, pp. 345–352.

Cason, Timothy N. and Daniel Friedman (1996). “Price Formation in Double Auction Markets”. In: *Journal of Economic Dynamics and Control* 20.8, pp. 1307–1337.

Cetorelli, Nicola, Beverly Hirtle, Donald P. Morgan, Stavros Peristiani, and João Santos (2007). “Trends in financial market concentration and their implications for market stability”. In: *Economic Policy Review* 13.1, pp. 33–51.

Chakraborty, Mithun, Sanmay Das, and Justin Peabody (2015). “Price Evolution in a Continuous Double Auction Prediction Market with a Scoring-Rule Based Market Maker”. In: *Twenty-Ninth AAAI Conference on Artificial Intelligence*. Austin, TX, pp. 835–841.

Chiarella, Carl and Giulia Iori (2002). “A simulation analysis of the microstructure of double auction markets”. In: *Quantitative Finance* 2.5, pp. 346–353.

Cliff, Dave (2009). “ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms”. In: *IEEE Transactions on Evolutionary Computation* 13, pp. 3–18.

Das, Rajarshi, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro (2001). “Agent-Human Interactions in the Continuous Double Auction”. In: *Seventeenth International Joint Conference on Artificial Intelligence*, pp. 1169–1178.

Daskalakis, Constantinos, Paul W. Goldberg, and Christos H. Papadimitriou (2009). “The complexity of computing a Nash equilibrium”. In: *SIAM Journal on Computing* 39.1, pp. 195–259.

Doob, J. L. (1942). “The Brownian Movement and Stochastic Equations”. In: *Annals of Mathematics* 43.2, pp. 351–369.

Du, Songzi and Haoxiang Zhu (2017). “What is the Optimal Trading Frequency in Financial Markets?” In: *Review of Economic Studies* 84.4, pp. 1606–1651.

Farmer, J. Doyne, Paolo Patelli, and Ilija I. Zovko (2005). “The Predictive Power of Zero Intelligence in Financial Markets”. In: *Proceedings of the National Academy of Sciences* 102, pp. 2254–2259.

Feng, Ling, Baowen Li, Boris Podobnik, Tobias Preis, and H. Eugene Stanley (2012). “Linking agent-based models and stochastic models of financial markets”. In: *Proceedings of the National Academy of Sciences* 109.22, pp. 8388–8393.

Fricke, Daniel and Austin Gerig (2018). “Too fast or too slow? Determining the optimal speed of financial markets”. In: *Quantitative Finance* 18.4, pp. 519–532.

Friedman, Daniel (1993). “The double auction market institution: A survey”. In: *The Double Auction Market: Institutions, Theories, and Evidence*. Ed. by Daniel Friedman and John Rust. Addison-Wesley, pp. 3–25.

Friedman, Daniel and John Rust, eds. (1993). *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley.

Fudenberg, Drew and Jean Tirole (1991). "Perfect Bayesian equilibrium and sequential equilibrium". In: *Journal of Economic Theory* 53.2, pp. 236–260.

Garbade, Kenneth D. and William L. Silber (1979). "Structural Organization of Secondary Markets: Clearing Frequency, Dealer Activity and Liquidity Risk". In: *Journal of Finance* 34.3, pp. 577–593.

Gjerstad, Steven and John Dickhaut (1998). "Price Formation in Double Auctions". In: *Games and Economic Behavior* 22.1, pp. 1–29.

Gode, Dhananjay K. and Shyam Sunder (1993). "Allocative Efficiency of Markets with Zero-intelligence Traders: Market as a Partial Substitute for Individual Rationality". In: *Journal of Political Economy*, pp. 119–137.

— (1997). "What makes markets allocationally efficient?" In: *Quarterly Journal of Economics* 112.2, pp. 603–630.

Goettler, Ronald L., Christine A. Parlour, and Uday Rajan (2009). "Informed Traders and Limit Order Markets". In: *Journal of Financial Economics* 93.1, pp. 67–87.

Golub, Anton, John Keane, and Ser-Huang Poon (2012). *High frequency trading and mini flash crashes*. Tech. rep. Available at SSRN, 2182097. University of Manchester.

Haas, Marlene and Marius A. Zoican (2016). "Beyond the Frequency Wall: Speed and Liquidity on Batch Auction Markets". In: *Paris Finance Meeting EUROFIDAI*.

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld (2011). "Does algorithmic trading improve liquidity?" In: *Journal of Finance* 66.1, pp. 1–33.

Hommes, Cars and Florian Wagener (2009). "Complex evolutionary systems in behavioral finance". In: *Handbook of financial markets: dynamics and evolution*. Ed. by Thorsten Hens and Klaus Reiner Schenk-Hoppé. North-Holland. Chap. 4, pp. 217–276.

How many stars are there in the universe? (2004). European Space Agency. URL: `http://www.esa.int/Our_Activities/Space_Science/Herschel/How_many_stars_are_there_in_the_Universe` (visited on 03/05/2017).

Johnson, Neil, Eric Zhao Guannan Hunsader, Hong Qi, Nicholas Johnson, Jing Meng, and Brian Tivnan (2013). "Abrupt rise of new machine ecology beyond human response time". In: *Scientific Reports* 3, p. 2627.

Jordan, Patrick R., L. Julian Schvartzman, and Michael P. Wellman (2010). "Strategy Exploration in Empirical Games". In: *9th International Conference on Autonomous Agents and Multiagent Systems*, pp. 1131–1138.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun (2017). "The Flash Crash: High-frequency trading in an electronic market". In: *Journal of Finance* 72.3, pp. 967–998.

Leal, Sandrine Jacob and Mauro Napoletano (2017). "Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading". In: *Journal of Economic Behavior and Organization*. ISSN: 0167-2681.

LeBaron, Blake (2006). "Agent-Based Computational Finance". In: *Handbook of Agent-Based Computational Economics*. Ed. by Leigh Tesfatsion and Kenneth L. Judd. Elsevier.

LeBaron, Blake and Ryuichi Yamamoto (2008). "The impact of imitation on long memory in an order-driven market". In: *Eastern Economic Journal* 34.4, pp. 504–517.

Lee, Bernard, Shih-Fen Cheng, and Annie Koh (2010). "An analysis of extreme price shocks and illiquidity among systematic trend followers". In: *Available at SSRN 1546458*.

— (2011). "Would Price Limits have Made any Difference to the 'Flash Crash' on May 6, 2010?" In: *Review of Futures Markets* 19, pp. 55–93.

Li, Zhuoshu and Sanmay Das (2016). "An Agent-Based Model of Competition between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs?" In: *15th International Conference on Autonomous Agents and Multi-Agent Systems*. Singapore, pp. 50–58.

Lux, Thomas and Michele Marchesi (1999). "Scaling and criticality in a stochastic multi-agent model of a financial market". In: *Nature* 397.6719, p. 498.

Madhavan, Ananth (1992). "Trading mechanisms in securities markets". In: *Journal of Finance* 47.2, pp. 607–641.

Menkveld, Albert J. and Bart Zhou Yueshen (2017). *The Flash Crash: A cautionary tale about highly fragmented markets*. Tech. rep. Available at SSRN, 2243520. VU University Amsterdam and INSEAD.

Meyer, Gregory, Nicole Bullock, and Joe Rennison (2018). *How high-frequency trading hit a speed bump*. Financial Times. (Visited on 03/05/2018).

Paddrik, Mark, Roy Hayes Jr., Andrew Todd, Steve Yang, Peter Beling, and William Scherer (2012). "An Agent Based Model of the E-Mini S&P 500 Applied to Flash Crash Analysis". In: *IEEE Conference on Computational Intelligence for Financial Engineering and Economics*, pp. 1–8.

Popper, Nathaniel and Jenny Anderson (2015). *Trader Arrested in Manipulation that Contributed to 2010 'Flash Crash'*. The New York Times. (Visited on 11/01/2016).

Rust, John, John Miller, and Richard Palmer (1993). "Behavior of Trading Automata in a Computerized Double Auction Market". In: *The Double Auction Market: Institutions, Theories, and Evidence*. Ed. by Daniel Friedman and John Rust. Addison-Wesley, pp. 155–198.

Satterthwaite, Mark A. and Steven R. Williams (1989). "Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency". In: *Journal of Economic Theory* 48.1, pp. 107–133.

Schnitzlein, Charles R. (1996). "Call and continuous trading mechanisms under asymmetric information: An experimental investigation". In: *Journal of Finance* 51.2, pp. 613–636.

Schvartzman, L. Julian and Michael P. Wellman (2009). "Stronger CDA strategies through empirical game-theoretic analysis and reinforcement learning". In: *8th International Conference on Autonomous Agents and Multiagent Systems*, pp. 249–256.

Schwartz, Robert A. and Liuren Wu (2013). "Equity Trading in the Fast Lane: The Staccato Alternative". In: *Journal of Portfolio Management* 39 (3), pp. 3–6.

SEC (2013). *Knight Capital Americas LLC*. Administrative Proceeding. U.S. Securities & Exchange Commission. URL: `https://www.sec.gov/litigation/admin/2013/34-70694.pdf`.

Sparrow, Chris (2012). "The Failure of Continuous Markets". In: *Journal of Trading* 7.2, pp. 44–47.

Tesauro, Gerald and Rajarshi Das (2001). "High-performance bidding agents for the continuous double auction". In: *Third ACM Conference on Electronic Commerce*. Tampa, pp. 206–209.

Vorobeychik, Yevgeniy, Christopher Kiekintveld, and Michael P. Wellman (2006). "Empirical Mechanism Design: Methods, with Application to a Supply Chain Scenario". In: *Seventh ACM Conference on Electronic Commerce*. Ann Arbor, MI, pp. 306–315.

Vorobeychik, Yevgeniy and Michael P. Wellman (2008). "Stochastic search methods for Nash equilibrium approximation in simulation-based games". In: *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 1055–1062.

Vytelingum, Perukrishnen, Dave Cliff, and Nicholas R. Jennings (2008). "Strategic bidding in continuous double auctions". In: *Artificial Intelligence* 172, pp. 1700–1729.

Wah, Elaine, Dylan R. Hurd, and Michael P. Wellman (2015). "Strategic Market Choice: Frequent call markets vs. continuous double auctions for fast and slow traders". In: *Conference on Auctions, Market Mechanisms and Their Applications*.

Wah, Elaine and Michael P. Wellman (2013). "Latency arbitrage, market fragmentation, and efficiency: A two-market model". In: *Fourteenth ACM Conference on Electronic Commerce*, pp. 855–872.

— (2016). "Latency Arbitrage in Fragmented Markets: A Strategic Agent-Based Analysis". In: *Algorithmic Finance* 5, pp. 69–93.

Wah, Elaine, Mason Wright, and Michael P. Wellman (2017). "Welfare Effects of Market Making in Continuous Double Auctions". In: *Journal of Artificial Intelligence Research* 59, pp. 613–650.

Wang, Xintong and Michael P. Wellman (2017). "Spoofing the Limit Order Book: An Agent-Based Model". In: *Sixteenth International Conference on Autonomous Agents and Multi-Agent Systems*. São Paulo, pp. 651–659.

Wellman, Michael P. (2006). "Methods for Empirical Game-theoretic Analysis (extended abstract)". In: *Twenty-First National Conference on Artificial Intelligence*, pp. 1552–1555.

— (2011). *Trading Agents*. Morgan & Claypool Publishers.

Wellman, Michael P., Tae Hyung Kim, and Quang Duong (2013). "Analyzing Incentives for Protocol Compliance in Complex Domains: A Case Study of Introduction-Based Routing". In: *Twelfth Workshop on the Economics of Information Security*.

Wiedenbeck, Bryce, Ben-Alexander Cassell, and Michael P. Wellman (2014). "Bootstrap statistics for empirical games". In: *International Conference on Autonomous Agents and Multi-Agent Systems*, pp. 597–604.

Wiedenbeck, Bryce and Michael P. Wellman (2012). "Scaling simulation-based game analysis through deviation-preserving reduction". In: *Eleventh International Conference on Autonomous Agents and Multi-Agent Systems*. Valencia, pp. 931–938.

Wiedenbeck, Bryce, Fengjun Yang, and Michael P. Wellman (2018). "A Regression Approach for Modeling Games with Many Symmetric Players". In: *32nd AAAI Conference on Artificial Intelligence*.

Wright, Mason and Michael P. Wellman (2018). "Evaluating the stability of non-adaptive trading in continuous double auctions". In: *Seventeenth International Conference on Autonomous Agents and Multi-Agent Systems*.

Zhan, Wenjie and Daniel Friedman (2007). "Markups in double auction markets". In: *Journal of Economic Dynamics and Control* 31, pp. 2984–3005.